arXiv:1511.02647v3 [cs.SI] 3 Jun 2016

Modelling influence and opinion evolution in online collective behaviour

Corentin Vande Kerckhove[1,+], Samuel Martin[2,+,*], Pascal Gend[2], Peter J. Rentfrow[3], Julien M. Hendrickx[1], Vincent D. Blondel[1]

1 Large graphs and networks group, Université catholique de Louvain, Avenue Georges Lemaitre, 4 B-1348 Louvain-la-Neuve, Belgium
2 Université de Lorraine, CRAN, UMR 7039 and CNRS, CRAN, UMR 7039, 2 Avenue de la Forêt de Haye, Vandoeuvre-les-Nancy, France
3 Department of Psychology, University of Cambridge

+ These authors contributed equally to this work.
*** samuel.martin@univ-lorraine.fr**

Abstract

Opinion evolution and judgment revision are mediated through social influence. Based on a large crowdsourced in vitro experiment ($n = 861$), it is shown how a consensus model can be used to predict opinion evolution in online collective behaviour. It is the first time the predictive power of a quantitative model of opinion dynamics is tested against a real dataset. Unlike previous research on the topic, the model was validated on data which did not serve to calibrate it. This avoids to favor more complex models over more simple ones and prevents overfitting. The model is parametrized by the *influenceability* of each individual, a factor representing to what extent individuals incorporate external judgments. The prediction accuracy depends on prior knowledge on the participants' past behaviour. Several situations reflecting data availability are compared. When the data is scarce, the data from previous participants is used to predict how a new participant will behave. Judgment revision includes unpredictable variations which limit the potential for prediction. A first measure of unpredictability is proposed. The measure is based on a specific control experiment. More than two thirds of the prediction errors are found to occur due to unpredictability of the human judgment revision process rather than to model imperfection.

Introduction

Many individual judgments are mediated by observing others' judgments. This is true for buying products, voting for a political party or choosing to donate blood. This is particularly noticeable on the online world. The availability of online data has lead to a recent surge in trying to understand how online social influence impact human behaviour. Some *in vivo* large scale online experiments were devoted to understand how information and behaviours spread in online social networks [1], others focused on determining which sociological attributes such as gender or age were involved in social influence processes [2].

Although decision outcomes are often tied to an objective best choice, outcomes can hardly be fully inferred from this supposedly best choice. For instance, predicting the popularity of songs in a cultural market requires more than just knowing the actual song quality [3]. The decision outcome is rather determined by the social influence process at work [4]. Hence, there is a need for opinion dynamics models with a predictive power.

Complementarily to the *in vivo* experiments, other recent studies used online *in vitro* experiment to identify the micro-level mechanisms susceptible to explain the way social influence

impacts human decision making [5–7]. These recent online *in vitro* studies have lead to posit that the so-called linear consensus model may be appropriate to describe the way individuals revise their judgment when exposed to judgments of others. The predictive power of such a mechanism remains to be assessed.

Trying to describe how individuals revise their judgment when subject to social influence has a long history in the psychological and social sciences. The consensus model used in this article draws from this line of work. These works were originally developed to better understand small group decision making. This occurs for instance when a jury in civil trials has to decide the amount of compensation awarded to plaintiffs [8–10]. Various types of tasks have been explored by researchers. These includes the forecasts of future events e.g., predicting market sales based on previous prices and other cues [11, 12], the price of products [13, 14], the probability of event occurrence [15, 16], such as the number of future cattle deaths [17], or regional temperatures [18]. The central ingredient entering in models of judgment revision is the weight which individuals put on the judgments of others, termed *influenceability* in the present article. This quantity is also known as the *advice taking weight* [14, 17] or the *weight of advice* [19–21]. It is represented by a number taking 0 value when the individual is not influenced and 1 when they entirely forget their own opinion to adopt the one from other individuals in the group. It has been observed that in a vast majority of the cases, the final judgment falls between the initial one and the ones from the rest of the group. Said otherwise, the influenceability lies between 0 and 1. This has been shown to sensibly improve the accuracy of decisions [22]. A 20% improvement has been found in an experiment when individuals considered the opinion of another person only [19]. However, individuals do not weight themselves and others equally. They rather overweight their own opinions [23]. This has been coined egocentric discounting [24]. Many factors affect influenceability. These include the perceived expertise of the adviser [17, 25] which may result from age, education, life experience [26], the difficulty of the task [21], whether the individual feels powerful [27] or angry [28], the size of the group [29], among others. A sensitivity analysis has been carried out to determine which factors most affect advice taking [30].

This line of work has focused on determining the factors impacting influenceability. None has yet answered whether judgment revision models could be used to predict future decisions. Instead, the models were validated on the data which served to calibrate the models themselves. This pitfall tends to favor more complex models over more simple ones and may result in overfitting. The model would then be unable to predict judgment revision from a new dataset. One reason for this literature gap could be the lack of access to large judgment revision database at the time, now made more readily available via online *in vitro* experiments. The predictability assessment is a necessary step to grow confidence in our understanding and in turn use this mechanism as a building block to design efficient online social systems. Revising judgments after being exposed to others' judgments takes an important role in many online social systems such as recommendation system [31, 32] or viral marketing campaign [33] among others. Unlike previous research, the present work provides an assessment of the model predictive power through crossvalidation of the proposed judgment revision model.

The prediction accuracy of a model is limited to the extent the judgment revision process is a deterministic process. However, there is theoretical [34–36] and empirical [37] evidence showing that the opinion individuals display is a sample of an internal probabilistic distribution. For instance, Vul and Pashler [37] showed that when participants were asked to provide their opinion twice with some delay in between, participants provided two different answers. Following these results, the present article details a new methodology to estimate the unpredictability level of the judgment revision mechanism. This quantifies the highest prediction accuracy one can expect.

The results presented in this article were derived using *in vitro* online experiments, where each participant repeated several times estimation tasks in very similar conditions. These repeated experiments yielded two complementary sets of results. First, it is shown that, in presence of social influence, the way individuals revise their judgment can be modeled using a quantitative model. Unlike the previously discussed studies, the gathered data allow assessing the predictive power of the model. The model casts individuals' behaviours according to their *influenceability*, the factor quantifying to what extent one takes external opinions into account. Secondly, a measure of intrinsic unpredictability in judgment revision is provided. Estimating the intrinsic unpredictability provides a limit beyond which no one can expect to improve

predictions. This last result was made possible through a specific *in vitro* control experiment. Although models of opinion dynamics have been widely studied for decades by sociologist [38] from a theoretical standpoint, to the best of our knowledge, it is the first time the predictive power of a quantitative model of opinion dynamics is tested against a real dataset.

Results and Discussion

To quantify opinion dynamics subject to social influence, we carried out online experiments in which participants had to estimate some quantities while receiving information regarding opinions from other participants. In a first round, a participant expresses their opinion x_i corresponding to their estimation related to the task. In the two subsequent rounds, the participant is exposed to a set of opinions x_j of other participants who performed the same task independently, and gets to update their own opinion. The objective of the study is to model and predict how an individual revises their judgment when exposed to other opinions. Two types of games were designed : the *gauging game*, in which the participants evaluated color proportions and the *counting game*, where the task required to guess amounts of items displayed in a picture (see *Experiment* section in *Material and Methods*). Participants to this online crowdsourced study were involved in 3-round judgment revision games. Judgment revision is modeled using a time-varying influenceability consensus model. In mathematical terms, $x_i(r)$ denotes the opinion of individual i at round r and its evolution is described as

$$x_i(r+1) \;=\; x_i(r) \;+\; \alpha_i(r) \cdot (\bar{x}(r) - x_i(r)) \tag{1}$$

where $r = 1, 2$ and where $\bar{x}(r)$ is the mean opinion of the group at round r (see *Opinion revision model* section in *Material and Methods* for details and supplementary section S3 for a test of the validity of the linearity assumption). This model is based on the *influenceability* $\alpha_i(r)$ of participants, a factor representing to what extent a participant incorporates external judgments.

Influenceability of participants

The influenceability is described for each participant by two parameters : $\alpha_i(1)$, the influenceability after first social influence and $\alpha_i(2)$, the influenceability after second social influence.

The distribution of couples $(\alpha_i(1), \alpha_i(2))$ were obtained by fitting model (1) to the whole dataset via mean square minimization for each type of games independently. The marginal distributions of $(\alpha_i(1), \alpha_i(2))$ are shown in Fig. 1. Most values fall within interval $[0, 1]$, meaning that the next judgment falls between one's initial judgment and the group judgment mean. Such a positive influenceability has been shown to improve judgment accuracy [22] (see also the *Practical Implications of the Model* section). Most individuals overweight their own opinion compared to the mean opinion to revise their judgment with $\alpha_i(r) < 0.5$. This fact is in accordance with the related literature on the subject [23].

An interesting research direction is to link the influenceability of an individual to their personality. One way to measure personality is via the big five factors of personality [39]. It turns out that influenceability is found not to be significantly correlated to the big five factors of personality, to education level and to gender. These negative findings are reported in the *Influenceability and personality* supplementary section.

The plots in Fig. 1 also display a small fraction of negative influenceabilities. One interpretation would be that the concerned participants recorded opinions very close to the average group opinion for multiple games. When this happens at one round, the opinion at the following round has a high probability to move away from the group average. This contributes to negative influenceabilities in the participants' influenceabilities during the fitting process.

Is the prediction error homogeneous for all influenceability values ? To see this, each bin of the influenceability distributions in Fig. 1 is colored to reflect average prediction error. The error is given in terms of root mean square error (*RMSE*). The color corresponds to the prediction error regarding participants having their influenceability falling within the bin. A detailed definition of RMSE is provided in paragraph *Validation procedure* in *Material and Methods*

section. This information shows that the model makes the best predictions for participants with a small but non-negative influenceability. On the contrary, predictions for participants with a high influenceability are less accurate.

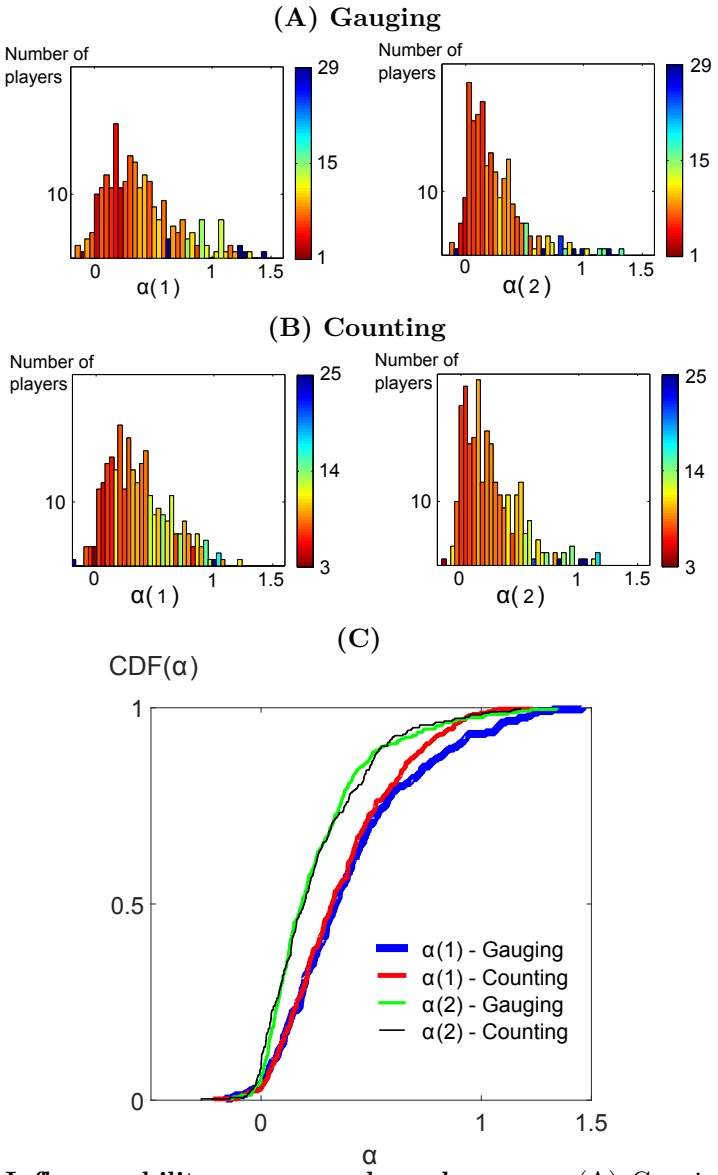

Figure 1. Influenceability over rounds and games. (A) Gauging game, (B) Counting game : distributions of the $\alpha_i(1)$ influenceability after the first round and $\alpha_i(2)$ influenceability after the second round for the time-varying influenceability model (1). The colormap corresponds to the average prediction $RMSE$ of participants in each bin. For visualization purposes, one value at $\alpha_i(1) = -1.5$ has been removed from the histogram in the gauging game, round 1. (There was only one individual with $\alpha_i(1) = -1.5$. For this particular individual, the linear relationship between $x_i(2) - x_i(1)$ and $\bar{x}(1) - x_i(1)$ is not significant (p-val = 0.38), so that the coefficient $\alpha_i(1) = -1.5$ should not be interpreted. For the rest of the participants, the level of trust can be read from the color scale. The color of a bar corresponds to the average prediction error made for the participants with α_i values falling within the bar range.) (C) Cumulated distributions of $\alpha_i(1)$ and $\alpha_i(2)$ for each type of games.

The distribution of influenceabilities of the population is subject to evolution over time. A two-sample Kolmogorov-Smirnov test rejects equality of distributions between $\alpha_i(1)$ and $\alpha_i(2)$ for both types of games (p-val $< 10^{-6}$, with KS distance of 0.25 and 0.23 for the *gauging game* and *counting game*, respectively). A contraction toward 0 occurs with a median going from 0.34 for $\alpha_i(1)$ to 0.18 for $\alpha_i(2)$ in the *counting game* and 0.32 to 0.20 in the *gauging game* (p-val$< 10^{-15}$ for a paired one-sided Wilcoxon sign-rank test). In other words, the participants continue to be influenced after round 2 but this influence is lightened. Fig. 1-C shows the discrepancy between the cumulated distribution functions over rounds.

Model Performance

Prediction scenarios When one wishes to predict how a participant revises their opinion in a decision making process, the level of prediction accuracy will highly depend on data availability. More prior knowledge on the participant should improve the predictions. When little prior information is available about the participant, the influenceability derived from it will be unreliable and may lead to poor predictions. In this case, it may be more efficient to proceed to a classification procedure provided that data from other participants are available. These approaches are tested by computing the prediction accuracy in several situations reflecting data availability scenarios.

In the worst case scenario, no data is available on the participant and the judgment revision mechanism is assumed to be unknown. In this case, predicting constant opinions over time is the only option. This corresponds to the null model against which the consensus model (1) is compared.

In a second scenario, prior data from the same participant is available. The consensus model can then be fitted to the data (individual influenceability method). Data ranging from 1 to 15 prior instances of the judgment process are respectively used to learn how the participant revises their opinion. Predictions are assessed in each of these cases to test how the predictions are impacted by the amount of prior data available.

In a final scenario, besides having access to prior data from the participant, it is assumed that a large body of participants took part in a comparable judgment making process. These additional data are expected to reveal the most common behaviours in the population and enable to derive typical influenceabilities by classification tools (population influenceability methods). For the population influenceability methods, there are two possibilities. First, assume that prior information on the participant is available. In this case, the influenceability class of the participant is determined using this information. In the alternative case, no prior data is available on the targeted participant. It is then impossible to discriminate which influenceability class they belong to. Instead, the most typical influenceability is computed for the entire population and the participant is predicted to follow this most typical behaviour.

Prediction accuracy We assess the predictions when the number of training data is reduced, accounting for realistic settings where prior knowledge on individuals is scarce. Individual parameter estimation via the individual influenceability method is compared to population influenceability method. This last method uses one or two $(\alpha(1), \alpha(2))$ couples of values derived on independent experiments. Fig. 2 presents the $RMSE$ (normalized to the range 0-100) obtained on validation sets for final round predictions using the model (1) with parameters fitted on training data of varying size. The methodology was also assessed for second round predictions instead of final round predictions. The results also hold in this alternative case, as described in *Second round predictions* section in *Material and Methods*.

The individual influenceability and population influenceability methods are compared to a null model assuming constant opinion with no influence, *i.e.,* $\alpha(r) = 0$ (*Null* in Fig. 2). The null model does not depend on training set size. By contrast, the individual influenceability method which, for each individual, fits parameters $\alpha_i(1)$ and $\alpha_i(2)$ based on training data, is sensitive to training set size (*Individual* α in Fig. 2) : it performs better than the null model when the number of games used for training is higher than 5 in both types of games but its predictions become poorer otherwise, due to overfitting.

Overfitting is alleviated using the population influenceability methods which restrict the choice of $\alpha(1)$ and $\alpha(2)$, making it robust to training size variations. The population method

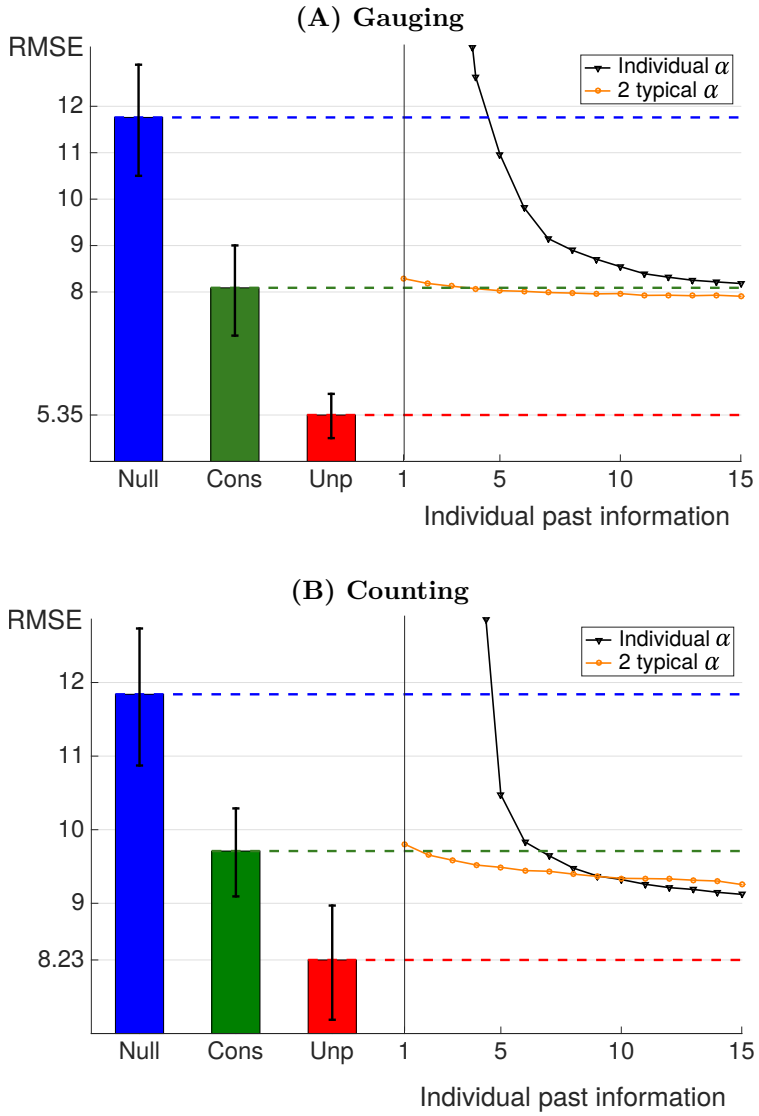

Figure 2. Root mean square error (*RMSE*) of the predictions for the final round. The RMSEs are obtained from crossvalidation. (A) *gauging game*, (B) *counting game*. In (B), the $RMSE$ has been scaled by a factor of 5 to be comparable to the (A) plot. The bar chart displays crossvalidation errors for models that does not depend on the training size. Top blue horizontal line corresponds to the null model of constant opinion. The middle horizontal green line corresponds to fitting using the same typical couple of influenceability for the whole population. The bottom horizontal red line corresponds to the prediction error due to intrinsic variations in judgment revision. The decreasing black curve (triangle dots) corresponds to fitting with the individual influenceability method. The slightly decreasing orange curve (round dots) corresponds to fitting choosing among 2 typical couples of influenceability. All $RMSE$ were obtained on validation games. The error bars provide 95% confidence intervals for the RMSEs.

which uses only one typical couple of influenceability as predictor presents one important advantage. It provides a method which does not require any prior knowledge about the participant targeted for prediction. It is thus insensitive to training set size (*Cons* in Fig. 2). This method improves by 31% and 18% the prediction error for the two types of games compared

to the null model of constant opinion.

The population methods based on two or more typical couples of influenceability require to possess at least one previous game by the participant to calibrate the model (*2 typical* α in Fig. 2). These methods are more powerful than the former if enough data is available regarding the participant's past behaviour (2 or 3 previous games depending on the type of games). The number of typical couples of influenceabilities to use depends on the data availability regarding the targeted participant. This is illustrated in Fig. 3. The modification obtained using more typical influenceabilities for calibration is mild. Moreover, too many typical influenceabilities may lead to poorer predictions due to overfitting. This threshold is reached for 4 couples of influenceabilities in the *gauging game* data. As a consequence, it is advisable to restrict the choice to 2 or 3 couples of influenceabilities. This analysis shows that possessing data from previous participants in a similar task is often critical to obtain robust predictions on judgment revision of a new participant.

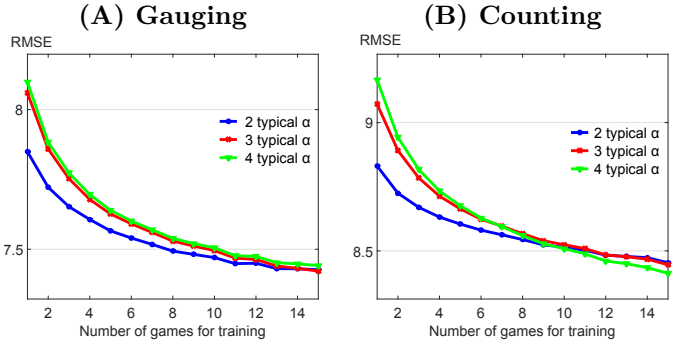

Figure 3. Predictive power of the consensus model when the number of typical couples of influenceability used for model calibration varies. Root mean square error (*RMSE*) plotted for training set size from 1 to 15 games. (A) *gauging game*, (B) *counting game*. In (B), *RMSE* has been scaled by a factor of 5 to be comparable to the (A) plot.

The results of the control experiments are displayed by a red dashed line in Fig. 2-A,B. This bottom line corresponds to the amount of prediction error which is due to the intrinsic unpredictability of judgment revision. No model can make better predictions than this threshold (see *Control experiment* section in *Material and Methods*).

The *gauging game* obtains an unpredictable *RMSE* of 5.35 while the *counting game* obtains 8.23. By contrast, the average square variation of the judgments between first and final rounds are respectively 11.76 and 11.84 for both types of games (corresponding to the *RMSE* of the null model). Taking the intrinsic unpredictable variation thresholds as a reference, the relative prediction *RMSE* is more than halved when using the time varying influenceability model (1) with one couple of typical influenceabilities instead of the null model with constant opinion. In other words, more than two thirds of the prediction error made by the consensus model is due to the intrinsic unpredictability of the decision revision process.

The error bars in Fig. 2 provide 95% confidence intervals for the RMSEs. They confirms statistical significance of the difference between RMSEs. For clarity, the error bars are provided only for regression methods which do not depend on training set size. For completeness, supplementary Fig. S3 provides error bars for all models.

RMSEs were used in this study since it corresponds to the quantity being minimized when computing the influenceability parameter α. Alternatively, reporting the *Mean Absolute Errors* (MAEs) may help the reader to obtain more intuition on the level of prediction error. For this reason, MAEs are provided in supplementary Fig. S4.

Practical Implications of the Model

Do groups reach consensus ? Because of social influence, groups tend to reduce their disagreement. However, this does not necessarily implies that groups reach consensus. To test how much disagreement remains after the social process, the distance between individual

judgments and mean judgments in corresponding groups is computed at each rounds. The results are presented for the *gauging game*. The same conclusions also hold for the *counting game*. Fig. 4 presents the statistics summary of these distances. The median distances are respectively 5.5, 4.2 and 3.5 for the three successive rounds, leading to a median distance reduction of 24% from round 1 to 2 and 16% from round 2 to 3. In other words, the contraction of opinion diversity is less important between rounds 2 and 3 than between rounds 1 and 2 and more than 50% of the initial opinion diversity is preserved at round 3. This is in accordance to the influenceability decay observed in *Influenceability of participants* section.

If one goes a step further and assumes that the contraction continues to lessen at the same rate over rounds, it may be that groups will never reach consensus. This phenomenon is quite remarkable since it would explain the absence of consensus without requiring non-linearities in social influence. An experiment involving an important number of rounds would shed light on this question and is left for future work.

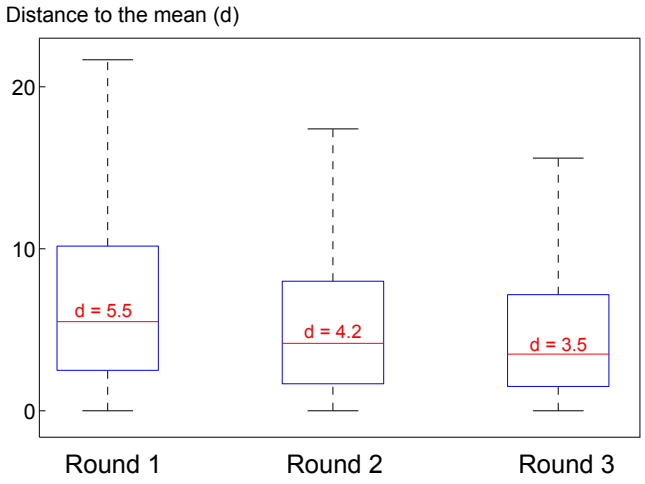

Figure 4. Distance of participants' judgment to mean judgment for rounds 1, 2 and 3 in the *gauging game*. Is displayed only the data coming from games where for all 3 rounds at least 5 out of 6 participants provided a judgment. This ensures that the judgment mean and standard deviation can be compared over rounds.

Influenceability and individual performance Each game is characterized by a true value, corresponding to an exact proportion to guess (for *gauging games*) or an exact amount of items displayed to the participants (for *counting games*). Whether social influence promotes or undermines individual performance can be measured for the two tasks. Individual performance can also be compared to the performance of the mean opinions in each group of 6 participants.

At each round r, a participant's success is characterized by the root mean square distance to truth, denoted $E_i(r)$. The error $E_i(r)$ depicts how far a participant is to truth. Errors are normalized to fit in the range $0-100$ in both types of tasks so as to be comparable. A global measure of individual errors is defined as the median over participants of $E_i(r)$, and the success variation between two rounds is given by the median value of the differences $E_i(r) - E_i(r+1)$. A positive or negative success variation corresponds respectively to a success improvement or decline of the participants after social interaction. The errors are displayed in Fig. 5. The results are first reported for the *gauging game*. The median error $E_i(r)$ for rounds 1, 2 and 3 are respectively 11.9, 10.0 and 9.8 (Fig. 5-(A)). It reveals an improvement with a success variation of 1.9 and 0.2 for $E_i(1) - E_i(2)$ and $E_i(2) - E_i(3)$ respectively (p-values$< 10E-5$, sign test), showing that most of the improvement is made between first and second round. Regarding the *counting game*, the median error $E_i(r)$ for rounds 1, 2 and 3 are respectively 23.1, 22.1 and 21.6 (Fig. 5-(B)). Note that the errors for the *counting game* have been rescaled by a factor of 5 to fit in the range $0-100$. This corresponds to an improvement with a success

variation of 1.0 and 0.5 for $E_i(1) - E_i(2)$ and $E_i(2) - E_i(3)$ respectively (the significance of the improvement is confirmed by a sign-test with p-values $< 10E-5$). Fig. 5 also reports the aggregate performance in terms of the root mean square distance from mean opinions to truth in each group of 6 participants. Unlike individual performance, the median aggregate performance does not consistently improve. Regarding the *gauging game*, the median aggregate error is significantly higher in round 1 than in rounds 2 and 3 (p-val$< 10E-4$, sign test) and this difference is not significant between rounds 2 and 3 (p-val> 0.05). Regarding the *counting game*, no significant difference is found among the 3 rounds for the median aggregate error (p-val> 0.05). As a consequence, social influence consistently helps the individual performance but does not consistently promote the aggregate performance.

The reason why social influence helps the individual performance is a combination of two factors. First, at round 1, the mean opinion is closer to the truth than the individual opinion is (p-val< 0.01, Mann–Whitney–Wilcoxon test). Second, in accordance to the consensus model (1), individuals move closer to the mean opinion over subsequent rounds. The fact that initially the mean opinion is closer to truth than the individual opinions corresponds to the wisdom of the crowd effect. The wisdom of the crowd is a statistical effect stating that averaging over several *independent* judgments yields a more accurate evaluation than most of the individual judgments would (see the early ox experiment by Galton in 1907 [40] or more recent work [41]). Since the aggregate performance does not consistently improve over rounds, it can be said that social influence does not consistently promote the wisdom of the crowd. Lorenz et al. [4] say that social influence *undermines* the wisdom of the crowd because it "reduces the diversity of the group without improving its accuracy". This variance reduction is also observed in the present study and corroborates the consensus model (1). Interestingly, in the first round, the wisdom of the crowd effect is more prominent in the *gauging game* than in the *counting game*: the median individual error is 37% higher than the median error of the mean opinion in the *gauging game* while it is only 8% higher in the *counting game*. The reason for this difference is studied in details in supplementary section S2.

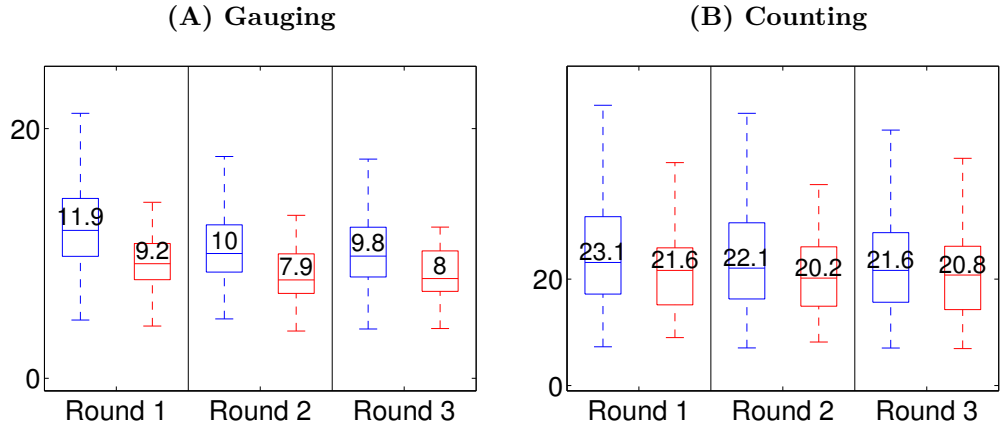

Figure 5. (A) *gauging game*; (B) *counting game*. From left to right : (Boxplot 1,3 and 5, in blue) Root mean square distance from individual opinions to truth ($E_i(r)$); (Boxplot 2,4 and 6 in red) Root mean square distance from mean opinion to truth. Values shown for the *counting game* have been scaled by a factor of 5.

The fact that the mean opinion is more accurate than individual opinions leads to posit that participants using the mean opinion to form their own opinion, *i.e.,* those with higher influenceability α_i, will increase their performance. We examine relationships between success variation and the model parameters α_i by computing partial Pearson correlations ρ controlling for the effect of the rest of the variables. Only significant Pearson correlations are mentioned (p-val < 0.05). All corresponding p-values happen to be smaller than 0.001 except for one as explicitly mentioned. Pearson correlations are given by pairs : the first value corresponds to

the *gauging game* while the second to the *counting game*. Influenceability $\alpha_i(1)$ between round 1 and 2 and improvement are positively related with $\rho(\alpha_i(1), E_i(1) - E_i(2)) = 0.41/0.22$. This is in accordance to the posited hypothesis. The *wisdom of the crowd effect* found at round 1 implies that participants who improve more from round 1 to 2 are those who give more weight to the average judgment. Since the wisdom of the crowd effect is more prominent in the *gauging game* than in the *counting game*, it is consistent that the correlation is higher in the former than in the latter. A similar effect relates success improvement and the influenceability $\alpha_i(2)$ between round 2 and 3 with $\rho(\alpha_i(2)\,,\, E_i(2) - E_i(3)) = 0.36/0.32$. As may be expected, higher initial success leaves less room for improvement in subsequent rounds, which explains that $\rho(E_i(1), E_i(1) - E_i(2)) = 0.68/0.21$ and $\rho(E_i(1), E_i(2) - E_i(3)) = 0.25/0.16$ (where 0.16 is significant for p-val < 0.01). This also means that initially better participants are not better than average at using external judgments.

Modelling influenceability across different types of games The assessment of the predictive power of model (1) on both types of games provides a generalisability test of the prediction method. The two types of games vary in difficulty. The root mean square relative distance $E_i(1)$ between a participant's first round judgment and truth is taken as the measure of inaccuracy for each participant. The median inaccuracy for the *counting game* is 23.1 while it is 11.8 for the *gauging game* (Mood's median test supports the rejection of equal median, $p = 0$). Moreover, a Q-Q plot shows that inaccuracy is more dispersed for the *counting game*, suggesting that estimating quantities is more difficult than gauging proportion of colors.

The accuracy of model (1) is compared for the two datasets in Fig. 2. Interestingly, the model prediction ranks remains largely unchanged for the two types of games. As depicted in Fig. 1-C, influenceability distributions do not vary significantly between the two games. A two-sample Kolmogorov-Smirnov test fails to reject the equality of distribution null hypothesis of equal median with $p > 0.65$ and a KS distance of 0.06 for both $\alpha_i(1)$ and $\alpha_i(2)$. This means that although the participants have an increased difficulty when facing the *counting game*, they do not significantly modify how much they take judgments from others into account. Additionally, the relationships between participants' success and influenceability are preserved for both types of games. The preserved tendencies corroborate the overall resemblance of behaviours across the two types of games. These similarities indicate that the model can be applied to various types of games with different level of difficulty.

Conclusions

The way online social systems are designed has an important effect on judgment outcome [42]. Operating or acting on these online social systems provides a way to significantly impact our markets, politics [43] and health. Understanding the social mechanisms underlying opinion revision is critical to plan successful interventions in social networks. It will help to promote the adoption of innovative behaviours (e.g., quit smocking [44], eat healthy) [45]. The design and validation of models of opinion revision will enable to create a bridge between system engineering and network science [46].

The present work shows that it is possible to model opinion evolution in the context of social influence in a predictive way. When the data regarding a new participant is available, parameters best representing their influenceability are derived using mean-square minimization. When the data is scarce, the data from previous participants is used to predict how the new participant will revise their judgments. To validate our method, results were compared for two types of games varying in difficulty. The model performs similarly in the two experiments, indicating that our influenceability model can be applied to other situations.

The decaying influenceability model after being fit to the data suggests that despite opinion settlement, consensus will not be reached within groups and disagreement will remain. This suggests that there needs to be incentives for a group to reach a consensus. The analysis also reveals that participants who improve more are those with highest influenceability, this independently of their initial success.

The degree to which one may successfully intervene on a social system is directly linked to the degree of predictability of opinion revision. Because there must always be factors which fall

out of the researcher's reach (changing mood or motivations of participants), part of the process cannot be predicted. The present study provides way to assess the level of unpredictability of an opinion revision mechanism. This assessment is based on a control experiment with hidden replicated tasks.

The proposed experiment type and validation method can in principle be generalized to any sort of continuous judgment revision. The consensus model can also serve as a building block to more complex models when collective judgments rely on additional information exchange.

Material and Methods

Experiment

Our research is based on an experimental website that we built, which received participants from a crowdsourcing platform. When a participant took part in an experiment, they joined a group of 6 participants. Their task was to successively play 30 games of the same sort related to 30 distinct pictures.

Criteria for online judgment revision game The games were designed to reveal how opinions evolve as a result of online social influence. Suitable games have to satisfy several constraints. First, to finely quantify influence, the games ought to allow the evolution of opinion to be gradual. Numbers were chosen as the way for participant to communicate their opinion. Multiple choice questions with a list of unordered items (e.g., choosing among a list of holiday locations) were discarded. Along the same lines, the evolution of opinion requires uncertainty and diversity of a sufficient magnitude in the initial judgments. The games were chosen to be sufficiently difficult to obtain this diversity. Thirdly, to encourage serious behaviours, the participants were rewarded based on their success in the games. This required the accuracy of a participant to be computable. Games were selected to have an ideal opinion or *truth* which served as a reference. Subjective questions involving for instance political or religious opinions were discarded.

Additionally, the game had to satisfy two other constraints related to the online context where, unlike face-to-face experiments, the researcher cannot control behavioural trustworthiness. Since the educational and cultural background of participants is a priori unknown, the game had to be *accessible*, i.e., any person which could read English had to be able to understand and complete the game. As a result, the games had to be as simple as possible. For instance, games could not involve high-level mathematical computations. Despite being simple to understand our games were still quite difficult to solve, in accordance with the first constraint. Lastly, to anticipate the temptation to cheat, the solution to the games had to be absent from the Internet. Therefore, questions such as estimating the population of a country were discarded.

Gauging and counting games Each game was associated with a picture. In the *gauging game*, the pictures were composed of 3 colors and participants estimated the percentage as a number between 0 and 100 of the same given color in the picture. In the *counting game*, the picture was composed of between 200 and 500 many small items, so that the participant could not count the items one by one. The participants had then to evaluate the total number of these items as a number between 0 and 500. A game was composed of 3 rounds. The picture was kept the same for all 3 rounds. In each round, the participant had to make a judgment. During the first round, each of the 6 participants provided their judgment, independently of the other participants. During the second round, each participant anonymously received all other judgments from the first round and provided their judgment again. The third round was a repetition of the second one. Accuracy of all judgments were converted to a monetary bonus to encourage participants to improve their judgment at each round. Screenshots of the games' interface are provided in the *Design of the Experiment* section.

Design of the experiment The present section describes the experiment interface. A freely accessible single player version of the games was also developed to provide a first hand experience of the games. In the single player version, participants are exposed to judgments

stored on our database obtained from real participant in previous games. The single player version is freely accessible at *http://collective-intelligence.cran.univ-lorraine.fr/eg/login*. The interface and the timing of the single player version is the same as the version used in the control experiment. The only difference is that the freely accessible version does not involve redundant games and provides accuracy feedback to the participants.

In the multi-player version which was used for the uncontrolled experiment, the participants came from the *CrowdFlower*® external crowdsourcing platform where they received the URL of the experiment login page along with a keycode to be able to login. The ad we posted on CrowdFlower was as follows :

> ***Estimation game regarding color features in images***
> *You will be making estimations about features in images. Beware that this game is a 6-player game. If not enough people access the game, you will not be able to start and get rewarded. To start the game : click on <estimation-game> and login using the following information :*
> *login : XXXXXXXX*
> *password: XXXXXXXX*
> *You will receive detailed instruction there. At the end of the game you will receive a reward code which you must enter below in order to get rewarded :*
> < >

The participants were told they will be given another keycode at the end of the experiment which they had to use to get rewarded on the crowdsourcing platform, this forced the participants to finish the experiment if they wanted to obtain a payment. Secondly, the participants arrived on the experiment login page, chose a login name and password so they could come back using the same login name if they wanted to for another experiment (see supplementary Fig. S5). Once they had logged in, they were requested to agree on a consent form mentioning the preservation of the anonymity of the data (see the *Consent and privacy* section below for details). Thirdly, the participants were taken to a questionnaire regarding personality, gender, highest level of education, and whether they were native English speaker or not (all the experiment was written in English). The questions regarding personality come from a piece of work by Gosling and Rentfrow [39] and were used to estimate the five general personality traits. The questionnaire page is reported in supplementary Fig. S6. Once the questionnaire submitted, the participants have access to the detailed instructions on the judgment process (supplementary supplementary Fig. S7). After this step, they were taken to a waiting room until 6 participants had arrived at this step. At this point, they started the series of 30 games which appeared 3 at a time, with one lone round where they had to make judgments alone and two social rounds where the provided judgment being aware of judgments from others. An instance of the lone round is given in supplementary Fig. S8-(A) for the *counting game* while a social round is shown in supplementary Fig. S9. Instances of pictures for the *gauging game* are provided in supplementary Fig. S8-(B). In the *gauging game*, the question was replaced by "*What percentage of the following color do you see in the image ?*". For this type of games, a sample of the color to be gauged for each pictures was displayed between the question and the 3 pictures. At the end of the 30 games, the participants had access to a debrief page where was given the final score and the corresponding bonus. They could also provide a feedback in a text box. They had to provide their email address if they wanted to obtain the bonus (see supplementary Fig. S10).

Consent and privacy Before starting the experiment, participants had to agree electronically on a consent form mentioning the preservation of the anonymity of the data :

> *Hello! Thank you for participating in this experiment. You will be making estimations about features in images. The closer your answers are to the correct answer, the higher reward you will receive. Your answers will be used for research on personality and behaviour in groups. We will keep complete anonymity of participants at all time. If you consent you will first be taken to a questionnaire. Then, you will get to a detailed instruction page you should read over before starting the game. Do you understand and consent to the terms of the experiment explained above ? If so click on I agree below.*

In this way, participants were aware that the data collected from their participation were to be be used for research on personality and behaviour in groups. IP addresses were collected. Email addresses were asked. Email addresses were only used to send participants bonuses via Paypal® according to their score in the experiments. IP addresses were used solely to obtained the country of origin of the participants. Behaviours were analyzed anonymously. Information collected on the participants were not used in any other way than the one presented in the manuscript and were not distributed to any third party. Personality, gender and country of origin presented no correlation with influenceability or any other quantity reported in the manuscript. The age of participants was not collected. Only adults are allowed to carry out microtasks on the CrowdFlower platform : CrowdFlower terms and conditions include : "you are at least 18 years of age". The experiment was declared to the Belgian Privacy Comission (https://www.privacycommission.be/) as requested by law. The French INSERM IRB read the consent procedure and confirmed that their approval was not required for this study since the data were analyzed anonymously.

Control experiment Human judgment is such a complex process that no model can take all its influencing factors into account. The precision of the predictions is limited by the intrinsic variation in the human judgment process. To represent this degree of unpredictability, we consider the variation in the judgment revision process that would occur if a participant were exposed to two replicated games in which the set of initial judgments happened to be identical. A control experiment served to measure this degree of unpredictability.

To create replicated experimental conditions, the judgments of five out of the six participants were synthetically designed. The only human participant in the group was not made aware of this, so they would act as in the uncontrolled experiments. Practically, participants took part in 30 games, among these, 20 games had been designed to form 10 pairs of replicated games with an identical picture used in both games of a pair. To make sure the participants did not notice the presence of replicates, the remaining 10 games were distributed between the replicates. The order of appearance of the games with replicates is as follows : $\mathbf{1}, \mathbf{2}, 11|\mathbf{3}, \mathbf{4}, 12|\mathbf{6}, 13, \mathbf{5}|14, \mathbf{7}, \mathbf{8}|\mathbf{9}, \mathbf{1}, 15|\mathbf{4}, \mathbf{10}, 16|\mathbf{2}, \mathbf{6}, 17|\mathbf{7}, 18, \mathbf{3}|\mathbf{8}, \mathbf{5}, 19|\mathbf{10}, \mathbf{9}, 20$, where games 1 to 10 are the replicated games. The games successively appeared three at a time from left to right. The 15 synthetic judgments (5 participants over 3 rounds) which appeared in the first instance of a pair of replicates were copies of past judgment made by real participants in past uncontrolled experiments. The copied games collected in uncontrolled experiments were randomly selected among the games in which more than 5 participants had provided judgments. Since the initial judgment of the real participant could not be controlled, the 15 synthetic judgments in the second replicate had to be shifted in order to maintain constant the initial judgment distances in each replicate. The shift was computed in real time to match the variation of the real participant initial judgments between the two replicates. The same shift was applied to all rounds to keep the synthetic judgments consistent over rounds (see Fig. 6 for the illustration of the shifting process). This provided exactly the same set of initial judgments up to a constant shift in each pair of replicated games. Such an experimental setting allowed assessing the degree of unpredictability in judgment revision (see the *Prediction accuracy* section in *Results* for details).

Participants

Uncontrolled experiment The data were collected during July, September and October 2014. Overall, 654 distinct participants took part in the study (310 in the *gauging game* only, 308 in the *counting game* only and 36 in both). In total, 64 groups of 6 participants completed a *gauging game*, while 71 groups of 6 participants completed a *counting game*. According to their IP addresses, participants came from 70 distinct countries. Participants mostly originated from 3 continents : 1/3 from Asia, 1/2 from Europe and 1/6 from South America. As detailed at the end of the paragraph, most participants completed most of the 30 games and played trustworthfully. The others were ignored from the study via two systematic filters. First, since the prediction method was tested using up to 15 games in the model parameter estimation process, the predictions reported in present study concern only the participants who completed more than 15 out of the 30 games. This ensures that the number of games used for parameter

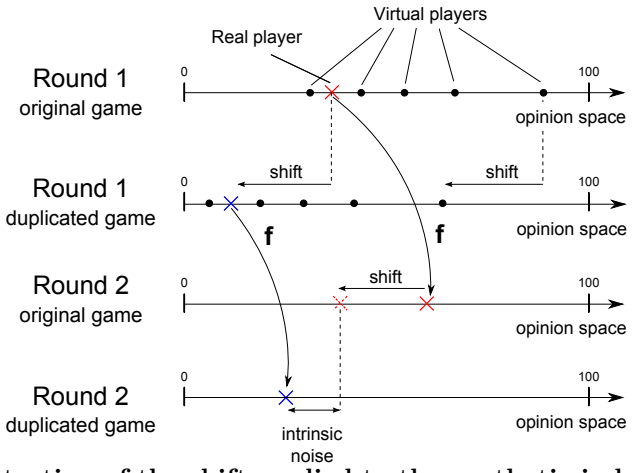

Figure 6. Illustration of the shift applied to the synthetic judgment in order to preserve the distances between initial judgments. The shift made by the participant in the first round was synthetically applied to all other initial judgments. The intrinsic variation in judgment revision is the distance between second round judgments, reduced by the shift in initial judgments. The dotted cross is not a judgment and is only displayed to show how the second round intrinsic variation is computed using the initial shift. All shifts are equal. The same shift is also applied to all other second round judgments, and used to compute final round intrinsic variations.

estimations is homogeneous over all participants. The prediction performance can then be compared among participants. The median number of fully completed games per participants was 24 with 10.5 std for the *gauging game* and 27 with 10.5 std for the *counting game*. Lower numbers are possibly due to loss of interest in the task or connexion issues. The first filter lead to keep 68% of the participants for the *gauging game* and 71% for the *counting game* (see Fig. 7–A,C for details). Secondly, the prediction were only made on judgments of trustworthy participants. Trustworthiness was computed via correlation between participant's judgments and true answers. Most participants carried out the task truthworthfully with a median correlation of 0.85 and median absolute deviation (MAD) of 0.09 for the *gauging game* and 0.70 median and 0.09 MAD for the *counting game*. A few participants either played randomly or systematically entered the same aberrant judgment. A minimum Pearson correlation thresholds of 0.61 for the *gauging game* and 0.24 for the *counting game* were determined using Iglewicz and Hoaglin method based on median absolute deviation [47]. The difference between the two thresholds is due to the higher difficulty of the *counting game* as expressed by the difference between median correlations. This lead to keep 91% and 96% of the participants which had passed the first filter (see Fig. 7–B,D for details).

It should be acknowledged that the selection procedure may have lead to sample bias. This could be due to self-selection : some people choose to participate in the experiment and others do not. The fact that the study was carried out online is another factor that could bias the sample. The *a posteriori* filters may be another source of bias. These are common issues in the behavioural sciences. Possibly, the nature of the study will have appealed to certain types of people and not others. Although that could have biased the characteristics of the sample, we are unaware of any empirical evidence suggesting that people who like participating in this kind of tasks are more or less susceptible to social influence.

Control experiment The data were collected during May and June 2015. Overall, 207 distinct participants took part in the study (113 in the *gauging game* only, 87 in the *counting game* only and 7 in both). The 120 *gauging game* participants took part in 139 independent games while the 94 *counting game* participants were involved in 99 independent games. Each independent game was completed by 5 synthetic participants to form groups of 6. The same filters as those used in the uncontrolled experiment were applied to the participants in the

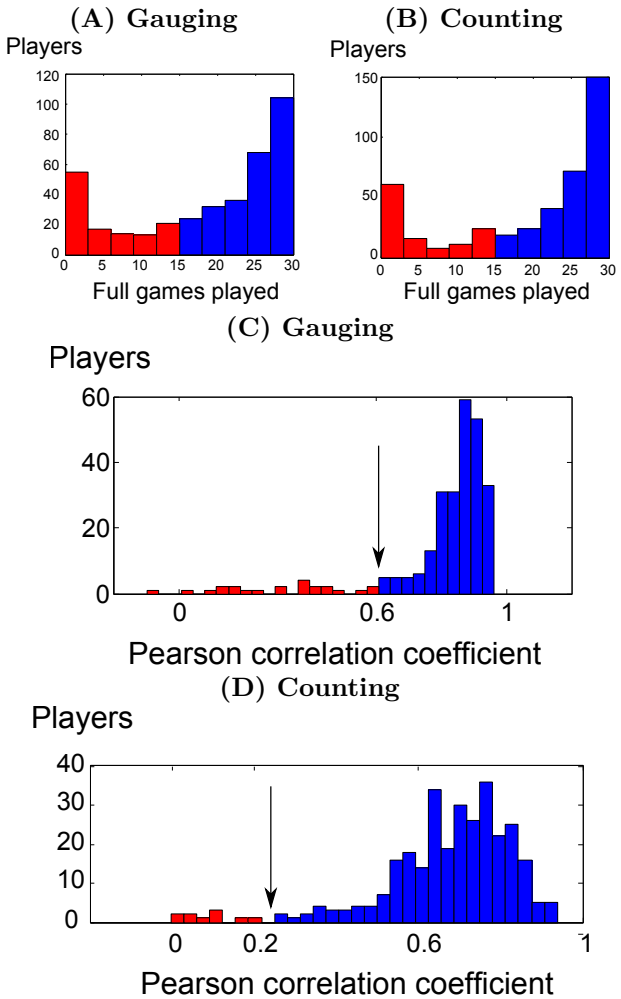

Figure 7. Filters for participants' trustworthiness. (A),(C) *gauging game*, (B),(D) *counting game*. (A,B) : Histograms of the number of full games played by participants. (C,D) : Histograms of the correlations between judgments and true values for each participant. Rejected participants are displayed in red while kept participants are in blue, to the right of the black arrow.

control experiment. This lead to keep 88% of the *counting games* and 80% of the *gauging games*.

Opinion revision model

To capture the way individual opinions evolve during a collective judgment process, a *consensus model* is used. These models have a long history in social science [38] and their behaviour has been thoroughly analyzed in a theoretical way [48, 49]. Our model (1) assumes that when an individual sees a set of opinions, their opinion changes linearly in the distance between their opinion and the mean of the group opinions. There is recent evidence supporting this assumption [7]. See also supplementary section S3 for a test of the validity of the linearity assumption. The rate of opinion change as a result of social influence is termed the *influenceability* of a participant. The model also assumes that this influenceability may vary over time. The decrease of influenceability represents opinion settling. The model is described in mathematical terms in equation (1), with $\alpha_i(r)$ being the influenceability of participant i after round r. When $\alpha_i(1)$ is nonzero, the ratio $\alpha_i(2)/\alpha_i(1)$ represents the decaying rate of influenceability. Parameters $\alpha_i(1)$ and $\alpha_i(2)$ are to be estimated to fit the model to the data. It is expected that

$\alpha_i(1), \alpha_i(2) \in [0, 1]$. If $\alpha_i(r) = 0$, the participant does not take into account the others and their opinion remains constant. If $\alpha_i(r+1) = \alpha_i(r)$, the influenceability does not change in time and the opinion $x_i(t)$ eventually converges to the mean opinion $\bar{x}$. Instead, if $\alpha_i(r+1)/\alpha_i(r) < 1$, influenceability starts positive but decays over time, which represents opinion settling.

There exist several variations to the linear consensus model presented above. In particular, the bounded confidence models [50,51] assume that the influenceability $\alpha_i(t)$ also depends on the distance between one's opinion and the influencing opinion. Alternatively, the model by Friedkin and Johnsen [52] assumes that individuals always remain influenced by their initial opinion or *prejudice* over time. Rather than providing an exhaustive assessment of the alternative models found in the literature, the objective of the present study is to show how the predictive power of a simple model of opinion dynamics can be assessed and to estimate the minimal prediction error that one can expect for any opinion dynamics model.

Consensus models of opinion dynamics are well adapted to represent opinions evolving in a continuous space. This corresponds to many real world situations in which the opinion represents the inclinaison between two extreme opposite options such as left and right in politics. Alternatively, part of the literature consider models with binary choices or actions (*e.g.,* voting for one candidates, buying a car or going on strikes). These discrete choice models include the rumour [53] and threshold models [54, 55]. In the latter, an individual changes their action when a certain proportion of its neighbours does so. These models directly link the discrete action of an individual to the actions in their neighbourhood. This allows to describe cascade propagation of behaviours. Presumably, before someone changes their actions, their opinion had to change as a result of the social stimuli. A recent model bridges these two bulks of work, considering the social influence of discrete actions on continuous opinions (the *CODA* model) which itself results in an individual action [56]. It appears that this model also naturally leads to cascade of behaviours over a social network [57]. It would be interesting to see how the threshold parameter in Granovetter's model may be expressed in terms of the initial opinion of individual in the CODA model : individuals with an opinion close to the boundary leading to either of the two discrete actions would have a lower threshold for action change.

The past two decades have witnessed a few attempts to confront simple models of opinion dynamics to real world data on collective human behaviours. These include conflicts and controversies on Wikipedia [58, 59] or how voters distribute their votes among candidates in elections [60–62]. However, since individual opinions involved in real world processes are not directly available, researchers had to calibrate their models on global measures such as level of controversy or distribution of votes. Moreover, the predictive power of these models was not assessed : the data used to calibrate the models also served to validate them. *In vitro* studies such as the present one have the advantage of providing the micro-level data driving the collective dynamics. Another advantage of *in vitro* studies is the possibility to differentiate social influence from other confounding factors such as homophily thanks to the anonymity of influencing individuals (see also [63, 64]).

Prediction procedure

Regression procedure The goal of the procedure is to predict the future judgment of a given participant in a given game. The set of first round judgments of this game are supposed to be known to initialize the model (1), this includes the initial judgment from the participant targeted for prediction and the initial judgments from the five other participants who possibly influenced the former.

To tune the model, prior games from the same participant are assumed to be available. These data serve to estimate the influenceability parameters $\alpha_i(1)$ and $\alpha_i(2)$. In one scenario, the influenceability parameters of the participant are estimated independently of the data from other participants (**individual influenceability** method). This is the only feasible method when no prior data on other participants is available (see the *Prediction scenarios* section in *Results* for details on the data availability scenarios). In this first case, parameters $\alpha_i(1)$ and $\alpha_i(2)$ are determined using a mean square error minimization procedure (this procedure amounts to likelihood maximization when the errors between model predictions and actual judgments are normally distributed and have zero mean, see for instance [65, p27]).

In situations where the number of prior games available to estimate the influenceability

parameters is small, little information is available on the participant's past behaviour. This may result in unreliable parameter estimates. To cope with such situations, another scenario is considered : besides having access to prior data from the targeted participant, part of the remaining participants (half of them, in our study) are used to derive the typical influenceabilities in the population. The expectation-maximization (EM) algorithm is used to classify the population into groups of similar influenceabilities [65]. These typical influenceabilities serve as a pool of candidates. The prior games of the targeted participant are used to determine which candidate yields the smallest mean square error (**population influenceability** method). The procedure to determine which typical candidate best suits a participant requires less knowledge than the one to accurately estimate their influenceability without a prior knowledge on its value (see the results in the *Prediction accuracy* section in *Results*).

Validation procedure The three scenarios presented in the *Prediction scenarios* section in *Results* are validated via crossvalidation. The validation procedure of the last scenario (i.e., access to prior data from participants, existence and access to typical influenceabilities) starts by randomly splitting the set of participants into two equal parts. The prediction focuses on one of the two halves while the remaining half is used to derive the typical influenceabilities in the population influenceability method. In the half serving for prediction, our model is assessed via repeated random sub-sampling crossvalidation : for each participant, a training subset of the games is used to assign the appropriate typical influenceabilities to each participant. The rest of the games serves as the validation set to compute the root mean square error ($RMSE$) between the observed data and predictions. The error specific to participant i is denoted as $RMSE_i$. The results are compared for various training set sizes. The $RMSE$ is obtained from averaging errors over 300 iterations using a different randomly selected training set each time. To compare scenarios 1 and 2 with scenario 3, we only consider the half serving for prediction. Scenario 1 (no data available) does not require any training step. For scenario 2 (access to prior data from participants), instead of learning affectation of typical influenceabilities to each participant, the learning process directly estimates the influenceabilities $\alpha_i(1)$ and $\alpha_i(2)$ for each participant without prior assumption on their values. The whole validation process is also carried out reversing the role of the two halves of the population and a global $RMSE$ is computed out of the entire process.

Intrinsic unpredictability estimation

Unpredictability of the second round judgment Even though the study mainly focuses on the predictions of third round judgments, we first focus in the present section on the two first rounds of the games, for the sake of clarity. The prediction procedure can easily be adapted to predict second round rather than third round judgments. Results for second round predictions based on the consensus model (1) are presented in *Second round predictions* section below.

The control experiment described in *Control experiment* section provides a way of estimating the intrinsic variations in the human judgment process. When a participant takes part in a game, their actual second judgment depends on several factors : their own initial judgment $x_i(1)$, the vector of initial judgments from other participants denoted as $x_{others}(1)$ and the displayed picture. As a consequence, the second round judgment of a participant can always be written as

$$x_i(2) = f_i^1(x_i(1), x_{others}(1), picture) + \eta, \tag{2}$$

where f_i^1 describes how a participant revises their judgment in average depending on their initial judgment and external factors. The term *picture* is the influence of the picture on the final judgment. The quantity η captures the intrinsic variation made by a participant when making their second round judgment despite having made the same initial judgment and being exposed to the same set of judgments and identical picture. Formally η is a random variable with zero mean. This is shown in supplementary section S1.1. The standard deviation of η is assumed to be the same for all participants to the same game, denoted as std(η). The standard deviation std(η) measures the root mean square error between f_i^1 and the actual judgment $x_i(2)$; this error measures by definition the intrinsic unpredictability of the judgment revision process. If it was known, the function f_i^1 would provide the best prediction regarding judgment

revision. By definition, no other model can be more precise. The function f_i^1 is unknown, but it is reasonable to make the following assumptions. First, the function f_i^1 is assumed to be a sum of (i) the influence of the initial judgments and and (ii) the influence of the picture. Thus f_i^1 splits into two components :

$$f_i^1(x_i(1), x_{others}(1), picture) =$$

$$\lambda g_i^1(x_i(1), x_{others}(1)) + (1-\lambda)h_i^1(picture),$$

where g_i^1 represents the dependence of the second round judgment on past judgments while h_i^1 contains the dependency regarding the picture. The parameter $\lambda \in [0,1]$ weights the relative importance of the first term compared to the second and is considered unique for each particular type of game. It is further assumed that if the initial judgment $x_i(1)$ and the others judgment at round 1 are shifted by a constant shift, the g_i^1 component in the second round judgment will on average be shifted in the same way, in other words, it is possible to write

$$g_i^1(x_i(1)+s, x_{others}(1)+s) = g_i^1(x_i(1), x_{others}(1)) + s,$$

where s is a constant shift applied to the judgments. Under this assumption, the control experiment provides a way of measuring the intrinsic variation. The intrinsic variation estimation $\mathrm{std}(\eta)$ can be empirically measured as the root mean of

$$\frac{1}{2}(x_i'(2) - x_i(2) - \lambda(x_i'(1) - x_i(1)))^2 \tag{3}$$

over all repeated games and all participants, where the prime notation is taken for judgments from the second replicated game in the control experiment (see *Participants* section in *Material and Methods*. The derivation of equation (3) is provided in supplementary section S4. Since η is assumed to have zero mean, the function f_i properly describing the actual judgment revision process is the one minimizing $\mathrm{std}(\eta)$. Correspondingly, the constant $\lambda \in [0,1]$ is set so as to satisfy this minimization. The intrinsic variation estimation $\mathrm{std}(\eta)$ is displayed in Fig. 8. The optimal λ values are found to be 0.67 and 0.74 and correspond to intrinsic unpredictability estimations of 5.38 and 7.36 for the *gauging game* and the *counting game*, respectively. These thresholds can be used to assess the quality of the predictions for the second round (see *Second round predictions* section).

Unpredictability of the final round judgment The procedure to estimate third round intrinsic unpredictability, whose results are shown in Fig. 2, varies slightly from second round estimation procedure. For third round judgments, a function with the same input as f_i^1, depending only on the initial judgments and the picture, cannot properly describe the judgment revision of one participant independently of the other participant's behaviour. In fact, the third round judgments depend on the second round judgments of others, which results from the revision process of other players. In other words, in a situation where a participant were to be faced successively to two groups of other participants, who by chance had given an identical set of initial judgments, the second round judgments of the two groups could vary due to distinct ways of revising their judgments.

Since the initial judgments does not suffice in describing third round judgments, function f_i^1 is modified to take the second round judgments of others $x_{others}(2)$ as an additional input. Then, the control experiment provides a way to estimate the intrinsic variations occurring in judgment revision up to the third round. This is described formally in the rest of the section. However, it should be noted that this description of judgment revision does not strictly speaking provide the exact degree of intrinsic variation included in the final round prediction error made in the uncontrolled experiment, it is rather an under-estimation of it : The predictions via the consensus model presented in the *Prediction performance* section in *Results*, are based solely on the initial judgments, whereas, the second round judgment from other participants are also provided in the present description. This additional piece of information necessarily makes the description of third round judgment more precise. As a consequence, the intrinsic variation estimated here (see details below) is an under-estimation of the actual intrinsic variation included in the prediction error of the consensus model. From a practical point of view, this means that the actual intrinsic unpredictability threshold is actually even closer to the predictions made

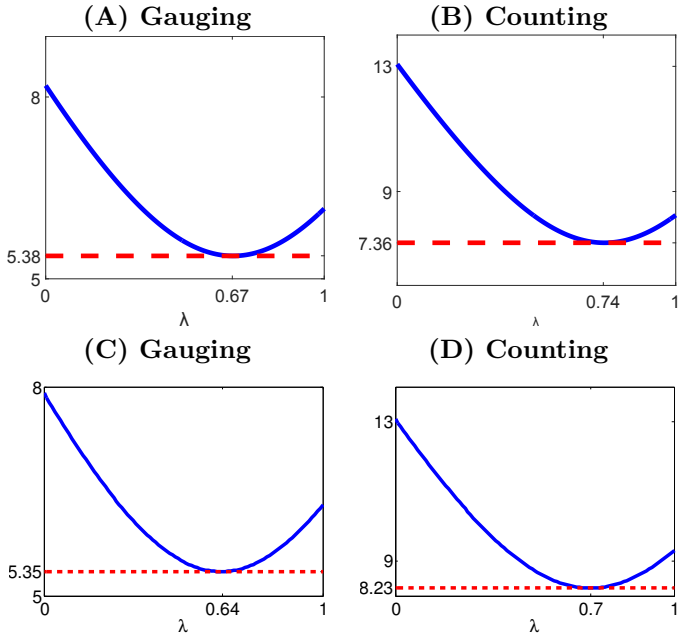

Figure 8. Root mean square intrinsic variation as a function of the λ parameters. (A), (B) : the second round judgments, as given in equation (3). (C), (D) : the third round judgments, as given in equation (5). (A), (C) *gauging game*, (B), (D) *counting game*.

by the consensus model than displayed in Fig. 2. In other words, there is even less space to improve the predictions provided by the consensus model, since more than two thirds of the error comes from the intrinsic variation rather than from the model imperfections.

Formally, the deterministic part of the third round judgment of a participant $x_i(3)$ is fully determined as a function of their initial judgment $x_i(1)$, the initial judgment of others $x_{others}(1)$, the second round judgment of others $x_{others}(2)$ and the picture. So, the third round judgment can be written as

$$x_i(3) = f_i^2(x_i(1), x_{others}(1), x_{others}(2), picture) + \bar{\eta}, \tag{4}$$

where $\bar{\eta}$ is the estimation of the intrinsic variation occurring after three rounds under the same initial judgments, same picture and same second round judgments from others. Under the same assumptions on function f_i^2 as those made on f_i^1, and analogously to equation (3), std($\bar{\eta}$) is measured by the root mean of

$$\frac{1}{2}(x_i'(3) - x_i(3) - \lambda(x_i'(1) - x_i(1)))^2 \tag{5}$$

over all repeated games and all participants. This intrinsic variation is provided as a function of parameter λ in Fig. 8, (C)-(D).

Second round predictions

The prediction procedure based on the consensus model (1) is applied to predict the second round judgments. Crossvalidation allows to assess the accuracy of the model. Results are presented in Fig. 9. These results are qualitatively equivalent to the prediction errors for the third round as shown in Fig. 2, although the second round predictions lead to a lower RMSEs, as expected since they correspond to shorter term predictions.

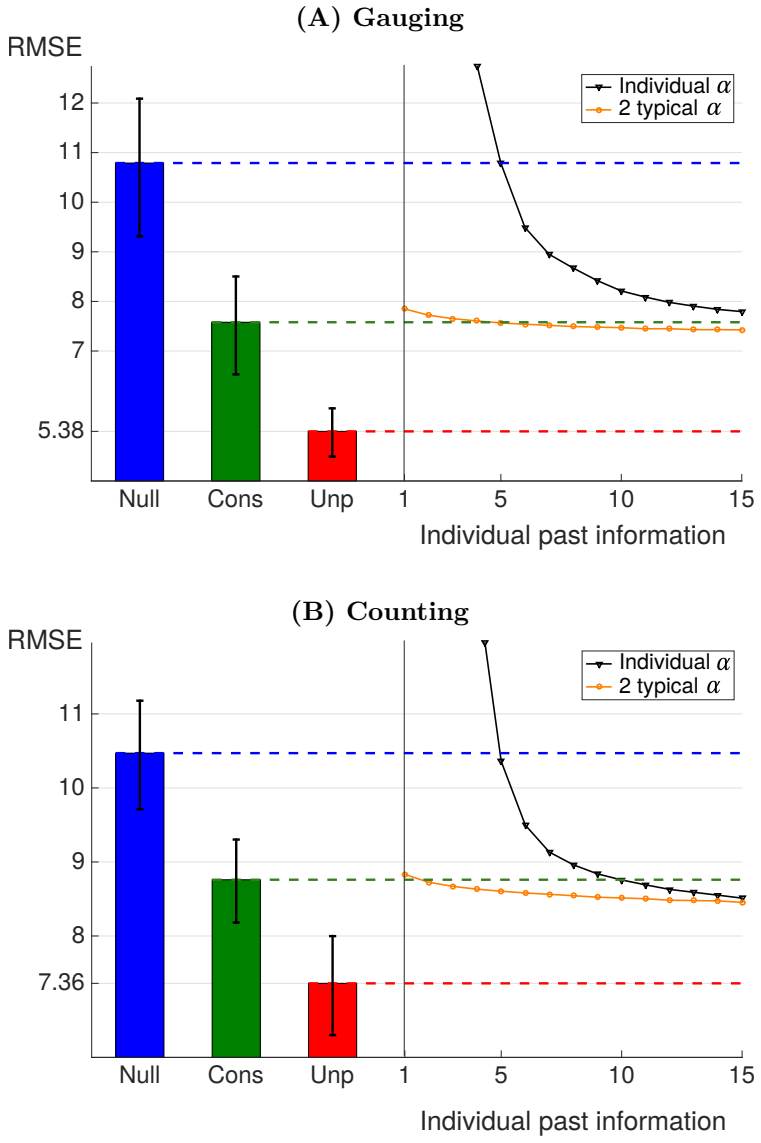

Figure 9. Root mean square error (*RMSE*) of the predictions for the *second* round. The RMSEs are obtained from crossvalidation. (A) *gauging game*, (B) *counting game*. In (B), the $RMSE$ has been scaled by a factor of 5 to be comparable to the (A) plot. The bar chart displays crossvalidation errors for models that does not depend on the training set size. Top blue horizontal line corresponds to the null model of constant opinion. The middle horizontal green line correspond to fitting using the same typical couple of influenceability for the whole population. The bottom horizontal red line correspond to the prediction error due to intrinsic variations in judgment revision. The decreasing black curves (triangle dots) correspond to fitting with the individual influenceability method. The slightly decreasing orange curves (round dots) correspond to fitting choosing among 2 typical couples of influenceability. All $RMSE$ were obtained on validation games.

Acknowledgments

C. V. K. is a F.N.R.S./FRIA research fellow. Research partly supported by the Belgian Interuniversity Attraction Poles (*Dynamical systems, control and optimization network*), by

an Actions incitatives 2014 of the Centre de Recherche en Automatique de Nancy, by project *Modélisation des dynamiques dans les réseaux d'échange de semences (PEPS MADRES)* funded by the Centre National pour la Recherche Scientifique and project *Computation Aware Control Systems (COMPACS)*, ANR-13-BS03-004 funded by the Agence National pour la Recherche. Contact : samuel.martin@univ-lorraine.fr. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author contributions statement

C. V. K., S. M., J. M. H., P. J. R. and V. D. B. conceived the experiment(s). C. V. K., S. M. and P. G. conducted the experiments. C. V. K. and S. M. analysed the results. All authors reviewed the manuscript.

Additional information

Competing financial interests: The authors declare no competing financial interests.

References

1. Centola D. The spread of behavior in an online social network experiment. science. 2010;329(5996):1194–1197.
2. Aral S, Walker D. Identifying influential and susceptible members of social networks. Science. 2012;337(6092):337–341.
3. Salganik MJ, Dodds PS, Watts DJ. Experimental study of inequality and unpredictability in an artificial cultural market. Science. 2006;311(5762):854–856.
4. Lorenz J, Rauhut H, Schweitzer F, Helbing D. How social influence can undermine the wisdom of crowd effect. Proceedings of the National Academy of Sciences. 2011;108(22):9020–9025.
5. Moussaïd M, Kämmer JE, Analytis PP, Neth H. Social influence and the collective dynamics of opinion formation. PloS one. 2013;8(11):e78433.
6. Chacoma A, Zanette DH. Opinion Formation by Social Influence: From Experiments to Modeling. PloS one. 2015;10(10):e0140406.
7. Mavrodiev P, Tessone CJ, Schweitzer F. Quantifying the effects of social influence. Scientific reports. 2013;3.
8. Hastie R, Penrod S, Pennington N. Inside the jury. The Lawbook Exchange, Ltd.; 1983.
9. Horowitz IA, ForsterLee L, Brolly I. Effects of trial complexity on decision making. Journal of applied psychology. 1996;81(6):757.
10. Hinsz VB, Indahl KE. Assimilation to Anchors for Damage Awards in a Mock Civil Trial1. Journal of Applied Social Psychology. 1995;25(11):991–1026.
11. Fischer I, Harvey N. Combining forecasts: What information do judges need to outperform the simple average? International journal of forecasting. 1999;15(3):227–246.
12. Harvey N, Harries C, Fischer I. Using advice and assessing its quality. Organizational behavior and human decision processes. 2000;81(2):252–273.
13. Schrah GE, Dalal RS, Sniezek JA. No decision-maker is an Island: integrating expert advice with information acquisition. Journal of Behavioral Decision Making. 2006;19(1):43–60.
14. Sniezek JA, Schrah GE, Dalal RS. Improving judgement with prepaid expert advice. Journal of Behavioral Decision Making. 2004;17(3):173–190.
15. Budescu DV, Rantilla AK. Confidence in aggregation of expert opinions. Acta psychologica. 2000;104(3):371–398.

16. Budescu DV, Rantilla AK, Yu HT, Karelitz TM. The effects of asymmetry among advisors on the aggregation of their opinions. Organizational Behavior and Human Decision Processes. 2003;90(1):178–194.

17. Harvey N, Fischer I. Taking advice: Accepting help, improving judgment, and sharing responsibility. Organizational Behavior and Human Decision Processes. 1997;70(2):117–133.

18. Harries C, Yaniv I, Harvey N. Combining advice: The weight of a dissenting opinion in the consensus. Journal of Behavioral Decision Making. 2004;17(5):333–348.

19. Yaniv I. The benefit of additional opinions. Current directions in psychological science. 2004;13(2):75–78.

20. Yaniv I. Receiving other people's advice: Influence and benefit. Organizational Behavior and Human Decision Processes. 2004;93(1):1–13.

21. Gino F. Do we listen to advice just because we paid for it? The impact of advice cost on its use. Organizational Behavior and Human Decision Processes. 2008;107(2):234–245.

22. Yaniv I, Milyavsky M. Using advice from multiple sources to revise and improve judgments. Organizational Behavior and Human Decision Processes. 2007;103(1):104–120.

23. Bonaccio S, Dalal RS. Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. Organizational Behavior and Human Decision Processes. 2006;101(2):127–151.

24. Yaniv I, Kleinberger E. Advice taking in decision making: Egocentric discounting and reputation formation. Organizational behavior and human decision processes. 2000;83(2):260–281.

25. Soll JB, Larrick RP. Strategies for revising judgment: How (and how well) people use others' opinions. Journal of Experimental Psychology: Learning, Memory, and Cognition. 2009;35(3):780.

26. Feng B, MacGeorge EL. Predicting receptiveness to advice: Characteristics of the problem, the advice-giver, and the recipient. Southern Communication Journal. 2006;71(1):67–85.

27. See KE, Morrison EW, Rothman NB, Soll JB. The detrimental effects of power on confidence, advice taking, and accuracy. Organizational Behavior and Human Decision Processes. 2011;116(2):272–285.

28. Gino F, Schweitzer ME. Blinded by anger or feeling the love: how emotions influence advice taking. Journal of Applied Psychology. 2008;93(5):1165.

29. Mannes AE. Are we wise about the wisdom of crowds? The use of group judgments in belief revision. Management Science. 2009;55(8):1267–1279.

30. Azen R, Budescu DV. The dominance analysis approach for comparing predictors in multiple regression. Psychological methods. 2003;8(2):129.

31. Pope DG. Reacting to rankings: evidence from "America's Best Hospitals". Journal of health economics. 2009;28(6):1154–1165.

32. Dellarocas C. The digitization of word of mouth: Promise and challenges of online feedback mechanisms. Management science. 2003;49(10):1407–1424.

33. Bessi A, Coletto M, Davidescu GA, Scala A, Caldarelli G, Quattrociocchi W. Science vs Conspiracy: collective narratives in the age of misinformation. PloS one. 2015;10(2):02.

34. Steyvers M, Griffiths TL, Dennis S. Probabilistic inference in human semantic memory. Trends in Cognitive Sciences. 2006;10(7):327–334.

35. Kersten D, Yuille A. Bayesian models of object perception. Current opinion in neurobiology. 2003;13(2):150–158.

36. Ma WJ, Beck JM, Latham PE, Pouget A. Bayesian inference with probabilistic population codes. Nature neuroscience. 2006;9(11):1432–1438.

37. Vul E, Pashler H. Measuring the crowd within probabilistic representations within individuals. Psychological Science. 2008;19(7):645–647.

38. French J. A formal theory of social power. Psychological Review. 1956;63:181–194.
39. Gosling SD, Rentfrow PJ, Swann WB. A very brief measure of the Big-Five personality domains. Journal of Research in personality. 2003;37(6):504–528.
40. Galton F. Vox populi (the wisdom of crowds). Nature. 1907;75:450–451.
41. Ariely D, Tung Au W, Bender RH, Budescu DV, Dietz CB, Gu H, et al. The effects of averaging subjective probability estimates between and within judges. Journal of Experimental Psychology: Applied. 2000;6(2):130.
42. Muchnik L, Aral S, Taylor SJ. Social influence bias: A randomized experiment. Science. 2013;341(6146):647–651.
43. Bond RM, Fariss CJ, Jones JJ, Kramer AD, Marlow C, Settle JE, et al. A 61-million-person experiment in social influence and political mobilization. Nature. 2012;489(7415):295–298.
44. Christakis NA, Fowler JH. The Collective Dynamics of Smoking in a Large Social Network. New England Journal of Medicine. 2008;358(21):2249–2258. PMID: 18499567.
45. Valente TW. Network interventions. Science. 2012;337(6090):49–53.
46. Liu YY, Slotine JJ, Barabási AL. Controllability of complex networks. Nature. 2011;473(7346):167–173.
47. Iglewicz B, Hoaglin DC. How to detect and handle outliers. vol. 16. ASQC Quality Press Milwaukee (Wisconsin); 1993.
48. Olfati-Saber R, Fax JA, Murray RM. Consensus and cooperation in networked multi-agent systems. Proceedings of the IEEE. 2007;95(1):215–233.
49. Martin S, Girard A. Continuous-time consensus under persistent connectivity and slow divergence of reciprocal interaction weights. SIAM Journal on Control and Optimization. 2013;51(3):2568–2584.
50. Deffuant G, Neau D, Amblard F, Weisbuch G. Mixing beliefs among interacting agents. Advances in Complex Systems. 2000;3(1–4):87–98.
51. Hegselmann R, Krause U. Opinion dynamics and bounded confidence models, analysis, and simulation. Journal of Artificial Societies and Social Simulation. 2002;5(3).
52. Friedkin NE, Johnsen EC. Social influence and opinions. Journal of Mathematical Sociology. 1990;15(3-4):193–206.
53. Dodds PS, Watts DJ. Universal behavior in a generalized model of contagion. Physical review letters. 2004;92(21):218701.
54. Granovetter M. Threshold models of collective behavior. American Journal of Sociology. 1978;83:1420–1443.
55. Watts DJ. A simple model of global cascades on random networks. Proceedings of the National Academy of Sciences. 2002;99(9):5766–5771.
56. Martins AC. Continuous opinions and discrete actions in opinion dynamics problems. International Journal of Modern Physics C. 2008;19(04):617–624.
57. Chowdhury N, Morarescu IC, Martin S, Srikant S. Continuous opinions and discrete actions in social networks: a multi-agent system approach. arXiv preprint arXiv:160202098. 2016;.
58. Iñiguez G, Török J, Yasseri T, Kaski K, Kertész J. Modeling social dynamics in a collaborative environment. EPJ Data Science. 2014;3(1):1–20.
59. Török J, Iñiguez G, Yasseri T, San Miguel M, Kaski K, Kertész J. Opinions, conflicts, and consensus: modeling social dynamics in a collaborative environment. Physical review letters. 2013;110(8):088701.
60. Bernardes AT, Stauffer D, Kertész J. Election results and the Sznajd model on Barabasi network. The European Physical Journal B-Condensed Matter and Complex Systems. 2002;25(1):123–127.

61. Caruso F, Castorina P. Opinion dynamics and decision of vote in bipolar political systems. International Journal of Modern Physics C. 2005;16(09):1473–1487.

62. Fortunato S, Castellano C. Scaling and universality in proportional elections. Physical Review Letters. 2007;99(13):138701.

63. Aral S, Muchnik L, Sundararajan A. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. Proceedings of the National Academy of Sciences. 2009;106(51):21544–21549.

64. Shalizi CR, Thomas AC. Homophily and contagion are generically confounded in observational social network studies. Sociological methods & research. 2011;40(2):211–239.

65. Bishop CM, et al. Pattern recognition and machine learning. vol. 1. Springer New York; 2006.

66. Niermann S. Testing for linearity in simple regression models. AStA Advances in Statistical Analysis. 2007;91(2):129–139.

67. Van der Linden D, te Nijenhuis J, Bakker AB. The general factor of personality: A meta-analysis of Big Five intercorrelations and a criterion-related validity study. Journal of research in personality. 2010;44(3):315–327.

Supplementary Information
Modelling influence and opinion evolution in online collective behaviour

Corentin Vande Kerckhove, Samuel Martin, Pascal Gend, Peter J. Rentfrow, Julien M. Hendrickx, Vincent D. Blondel

S1 Derivation of the measure of unpredictability

S1.1 Proof that η has zero mean

In section *Intrinsic unpredictability estimation*, we used that η has zero mean. This fact is proven in the sequel. Assume, to show a contradiction that $m = \mathbb{E}(\eta) \neq 0$. Denote $\tilde{x}_i(2) = f_i^1(x_i(1), x_{others}(1), picture)$. Then, we can show that $\tilde{x}_i(2) + m$ would be a better model for $x_i(2)$ than $\tilde{x}_i(2)$. This would contradict the definition of f_i^1 as the best model. Indeed, the expected square prediction error would be

$$\begin{aligned}\mathbb{E}\left((x_i(2) - (\tilde{x}_i(2) + m))^2\right) &= \mathbb{E}\left((x_i(2) - \tilde{x}_i(2))^2\right) - 2m\mathbb{E}\left(x_i(2) - \tilde{x}_i(2)\right) + m^2 \\ &= \mathbb{E}\left((x_i(2) - \tilde{x}_i(2))^2\right) - m^2 \\ &< \mathbb{E}\left((x_i(2) - \tilde{x}_i(2))^2\right),\end{aligned}$$

where we used the fact that $\eta = x_i(2) - \tilde{x}_i(2)$. The same reasoning allows to show that the prediction error $\bar{\eta}$ at round 3 also has zero mean.

S1.2 Derivation of equation (3)

Equation (3) is derived using the following reasoning. The judgments made in two replicated games of a control experiment by a same participants are described as

$$\begin{aligned}x_i(2) &= f_i^1(x_i(1), x_{others}(1), picture) + \eta, \\ x_i'(2) &= f_i^1(x_i'(1), x_{others}'(1), picture) + \eta',\end{aligned}$$

where the prime notation is taken for judgments from the second replicated game and η and η' are two independent draws of the random intrinsic variation. By design, the set of judgments are all shifted by the same constant :

$$\begin{aligned}x_i'(1) &= x_i(1) + s, \\ x_{others}'(1) &= x_{others}(1) + s,\end{aligned}$$

where $s = x_i'(1) - x_i(1)$ is known. According the assumption made on function f_i^1,

$$\begin{aligned}x_i(2) &= \lambda g_i^1(x_i(1), x_{others}(1)) + (1-\lambda)h_i^1(picture) + \eta, \\ x_i'(2) &= \lambda g_i^1(x_i'(1), x_{others}'(1)) + (1-\lambda)h_i^1(picture) + \eta',\end{aligned}$$

the second round judgment made in the second replicate is then

$$x_i'(2) = \lambda\left(g_i^1(x_i(1), x_{others}(1)) + s\right) + (1-\lambda)h_i^1(picture) + \eta',$$

where the invariance by translation of g_i^1 was used. Taking the difference makes the unknown terms $g_i^1(x_i(1), x_{others}(1))$ and $h_i^1(picture)$ vanish to obtain

$$x_i'(2) - x_i(2) = \lambda s + \eta' - \eta. \tag{6}$$

Since η and η' have zero mean and are assumed to have equal variance, the theoretical variance of η is

$$\begin{aligned}
\mathbb{E}(\eta^2) &= \frac{1}{2}\left(\mathbb{E}(\eta^2) + \mathbb{E}(\eta'^2)\right) \\
&= \frac{1}{2}(\mathbb{E}(\eta'^2) - 2\mathbb{E}(\eta'\eta) + \mathbb{E}(\eta^2) + 2\mathbb{E}(\eta'\eta)) \\
&= \frac{1}{2}(\mathbb{E}(\eta'^2 - 2\eta'\eta + \eta^2)) + \mathbb{E}(\eta'\eta) \\
&= \frac{1}{2}\mathbb{E}((\eta' - \eta)^2) + \mathbb{E}(\eta'\eta).
\end{aligned} \tag{7}$$

Moreover η and η' are assumed to be independent with zero mean, *i.e.,* , $\mathbb{E}(\eta) = \mathbb{E}(\eta') = 0$, therefore, their covariance is null : $\mathbb{E}(\eta\eta') = \mathbb{E}(\eta)\mathbb{E}(\eta') = 0$. Consequently, $\mathbb{E}(\eta^2) = \frac{1}{2}\mathbb{E}((\eta' - \eta)^2)$ and using equation (6), the variance of η is empirically measured as the average of

$$\frac{1}{2}(x_i'(2) - x_i(2) - \lambda s)^2$$

over all repeated games and all participants. This corresponds to equation (3).

S1.3 Discussion on the assumptions on η and η'

The only assumption used to derive equation (6) is that η and η' have the same variance and are independent for each participant. Since function f_i^1 is unknown, it is not possible to directly test these assumptions. However, since pairs of replicates in the control experiment are related to the same picture, it is unlikely that the covariance between η and η' would be negative. If the covariance was positive, the quantity given in equation (6) would become a lower bound on the unpredictability threshold, as shown through equation (7). Finally, if η and η' did not satisfy the assumption of equal variance, the quantity in equation (6) would still correspond to the average variance $\frac{1}{2}\left(\mathbb{E}(\eta^2) + \mathbb{E}(\eta'^2)\right)$ which also represents the average intrinsic unpredictability, as seen in equation (7).

S2 Circumstances of the wisdom of the crowd

The wisdom of the crowds may not always occur. The present section recalls one important hypothesis underlying the wisdom of the crowds. The hypothesis is then tested against the empirical data from the study. In the context of the present study, the wisdom of the crowd corresponds to the following fact: the mean opinion is most often *much* closer to the true answer than the individual opinions are. Denoting $\bar{x}$ the mean of n opinions x_i and T the corresponding true answer, this is formally expressed as

$$|\bar{x} - T| << \frac{1}{n}\sum_{i=1}^{n}|x_i - T|, \tag{8}$$

where $<<$ stands for *significantly smaller than*. The wisdom of the crowd given by equation 8 does not always take place. It only occurs if the opinions x_i are distributed sufficiently symmetrically around the true answer. When the distribution is largely biased above or below the true answer, equation 8 fails to

hold. To understand this fact, the group of individual is split in two : $i \in \mathcal{N}^+$ if $x_i(1) > T$ and $i \in \mathcal{N}^-$ if $x_i(1) < T$. Then, the distance of the mean opinion to truth rewrites as

$$\begin{aligned} |\bar{x} - T| &= \left| \left(\frac{1}{n} \sum_{i=1}^{n} x_i \right) - T \right| = \frac{1}{n} \left| \sum_{i=1}^{n} (x_i - T) \right| = \frac{1}{n} \left| \sum_{i \in \mathcal{N}^+} (x_i - T) + \sum_{i \in \mathcal{N}^-} (x_i - T) \right| \\ &= \frac{1}{n} \left| D^+ - D^- \right| \end{aligned}$$

where $D^+ = \sum_{i \in \mathcal{N}^+} |x_i - T| \geq 0$ is the contribution from opinions above truth and $D^- = \sum_{i \in \mathcal{N}^-} |x_i - T| \geq 0$ is the contribution from opinions below truth. Using these notation, the average distance to truth is $\frac{1}{n} \sum_{i=1}^{n} |x_i - T| = \frac{1}{n}(D^+ + D^-)$. As a consequence, the wisdom of the crowd described in equation 8 translates to

$$\left| D^+ - D^- \right| << (D^+ + D^-). \tag{9}$$

Two extreme cases are possible :

- **Perfect wisdom of the crowd :** opinions are homogeneously distributed around the true answer and $D^+ = D^-$ so that $|\bar{x}(1) - T| = 0$.
- **No wisdom of the crowd :** opinions either totally overestimate or totally underestimate the correct answer and either $D^- = 0$ or $D^+ = 0$, so that $|\bar{x} - T| = \frac{1}{n} \sum_{i=1}^{n} |x_i - T|$.

We now turn to the empirical data. Only the first round is discussed here because, in the subsequent rounds, the opinions are no longer independent, a criterion required for the wisdom of the crowd to occur. Fig. S1 displays how opinions are distributed around the true value for the *gauging game* (A) and the *counting game* (B). Both distributions fall between the two extreme cases with most opinions underestimating the true value. However, the bias is more important in the *counting game* which explains that the wisdom of the crowd is more prominent in the *gauging game* in the first round. This explains the differences between mean opinion errors and individual errors observed in Fig. 5.

S3 Testing the linearity of the consensus model

The consensus model (1) assumes that the opinion change $x_i(t+1) - x_i(t)$ grows linearly with the distance between $x_i(t)$ and the mean opinion $\bar{x}(t)$. This assumption is tested against the alternative

$$x_i(t+1) - x_i(t) = \beta_0 + \alpha_1 (\bar{x}(t) - x_i(t))^\gamma$$

with $\gamma \neq 1$. The numerical statistics values are reported for the opinion change between rounds 1 and 2 for the *gauging game*. The same conclusions hold for the *counting game* and for the opinion change between rounds 2 and 3. The linearity test provided in [66] applied to our data gives a statistics $P = -1.4\text{E}7$ with empirical variance $\text{var}(P) = 4\text{E}14$ so that we fail to reject the null hypothesis $\gamma = 1$ (p-val=0.5). Fig. S2 displays the evolution $x_i(t+1) - x_i(t)$ against the distance to the mean $\bar{x}(t) - x_i(t)$ along with the result of the linear regression assuming $\gamma = 1$.

S4 Influenceability and personality

Is influenceability related to personality ? To answer this question, we required the participants to provide information regarding their personality, gender, highest level of education, and whether they were native

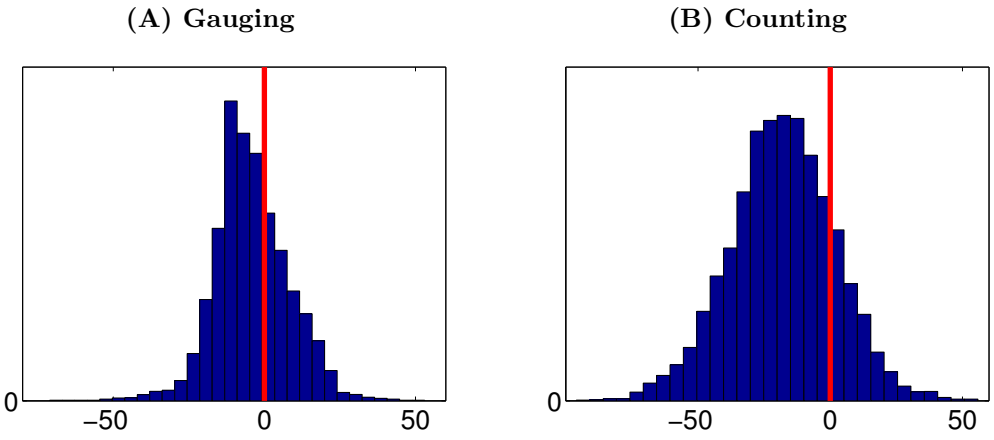

Figure S1. Deviation of opinions to truth during the lone round 1. Red vertical lines split the histogram into the opinions contributing negatively to the distance between mean opinion and truth (left) and the opinions contributing positively to the distance between mean opinion and truth (right). (A) *gauging game*; (B) *counting game*.

English speaker. The questionnaire regarding personality comes from a piece of work by Gosling and Rentfrow [39] and was used to estimate the five general personality traits. The questionnaire page is reported in supplementary Fig. S6. For each of the five traits, the participants rated how well they feel in adequacy with a set of synonyms (rating $s \in \{1, \ldots, 7\}$) and with a set of antonyms (rating $a \in \{1, \ldots, 7\}$. This redundancy allows for testing the consistency of the answer of each participants. The participants who had a distance $(8 - a) - s$ too far away from 0 were discarded (threshold values were found using Iglewicz and Hoaglin method based on median absolute deviation [47]). Partial Pearson's linear correlations are first reported between the individual traits measured by the questionnaire (see table S1). The correlation signs are found to be consistent with the related literature on the topic [67]. This indicates that our measure of the big five factors is trustworthy. Partial correlations are then provided to link the personal traits to influenceability. As shown in Table S2, none of the measured personal traits is able to explain the variability in the influenceability parameter. The only exceptions concern gender and being English native speaker, with weak level of significance (p-$val \in [0.01, 0.05]$). However, these relations are consistent neither between types of tasks nor over rounds, so that they cannot be trusted. We conclude that the big five personality factors and the other measured individual traits are not relevant to explain the influenceability parameter. Finding appropriate individual traits to explain the influenceability remains an open question.

(A) Participants from the *counting game*

	C	E	A	N	Gen	Eng	Edu
O	0.45***	0.38***	0.31***	-0.35***	0.03	-0.08	0.09
C		0.11	0.29***	-0.41***	-0.02	-0.06	0.29***
E			0.03	-0.15*	0.02	-0.05	0
A				-0.48***	0.11	-0.16*	0.15*
N					0.2**	0.08	-0.23***
Gen						0	0.01
Eng							-0.12

(B) Participants from the *gauging game*

	C	E	A	N	Gen	Eng	Edu
O	0.37***	0.24***	0.29***	-0.4***	0.06	0.15*	0.16**
C		0.2***	0.22***	-0.29***	0.15**	0.12*	0.11
E			-0.12*	-0.12*	0.04	0.04	-0.03
A				-0.34***	0.01	-0.07	0.08
N					0.18**	-0.02	-0.16**
Gen						-0.03	0
Eng							-0.02

Table S1. Partial Pearson's linear correlations among the big five factors of personality (O: openness, C : calmness, E : extroversion, A : agreeableness, N : neuroticism), gender (Gen), native English speaker (Eng) and highest level of education (Edu). Significance : *p-val< 0.05, **p-val< 0.01,***p-val< 0.001.

	(A) Gauging		**(B) Counting**	
	$\alpha(1)$	$\alpha(2)$	$\alpha(1)$	$\alpha(2)$
O	-0.1	-0.04	0.07	-0.05
C	0.08	-0.06	-0.05	-0.09
E	-0.1	0.02	-0.01	-0.01
A	0	-0.02	0.1	0.07
N	-0.06	0.03	-0.1	0.01
Gen	-0.14*	-0.05	0.03	0.12*
Eng	-0.09	0.02	-0.03	-0.12*
Edu	0.05	-0.1	-0.04	-0.1

Table S2. Partial Pearson's linear correlations linking influenceability to the big five factors of personality (O: openness, C : calmness, E : extroversion, A : agreeableness, N : neuroticism), gendre (Gen), native English speaker (Eng) and highest level of education (Edu). Significance : *p-val< 0.05, **p-val< 0.01,***p-val< 0.001.

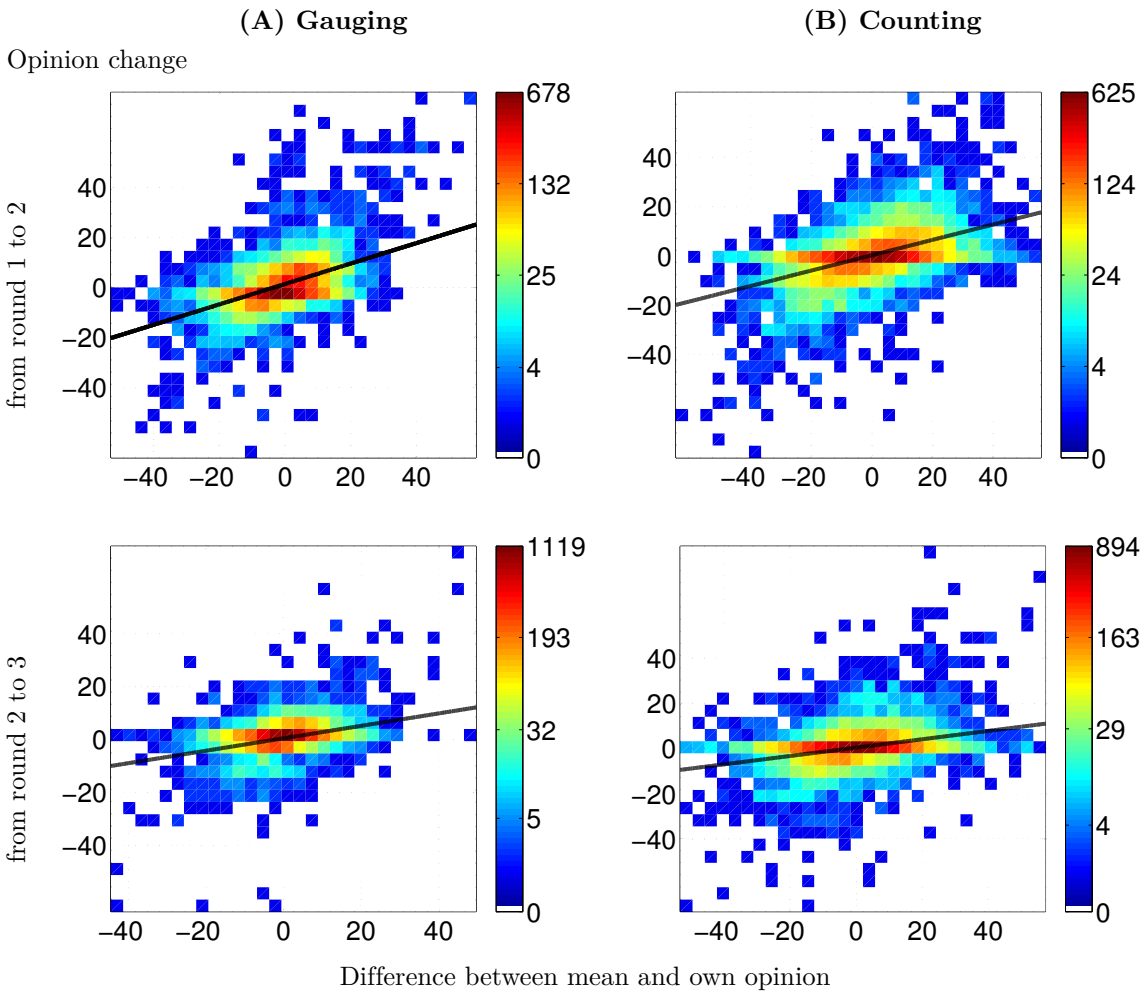

Figure S2. Opinion change $x_i(t+1) - x_i(t)$ versus difference between mean and individual opinion $\bar{x}(t) - x_i(t)$. The color in each cell corresponds to the number of data points falling in the cell. The color scale is logarithmic. The black strait line represents the linear regression $x_i(t+1) - x_i(t) = \beta_0 + \alpha(t)(\bar{x}(t) - x_i(t))$. Top : opinion change from round 1 to 2 ; bottom : opinion change from round 2 to 3. (A) *gauging game*; (B) *counting game*.

S5 Additional figures for prediction accuracy

S5.1 Confidence intervals for prediction errors

Fig. S3 displays error bars for 95% confidence interval of the RMSEs. This figure reveals that the two methods depending on training set size do not perform significantly better than the consensus model with one couple of typical influenceabilities, even for large training set sizes. This is an argument to favor the model in which the whole population has a unique couple of influenceabilities $(\alpha(1), \alpha(2))$.

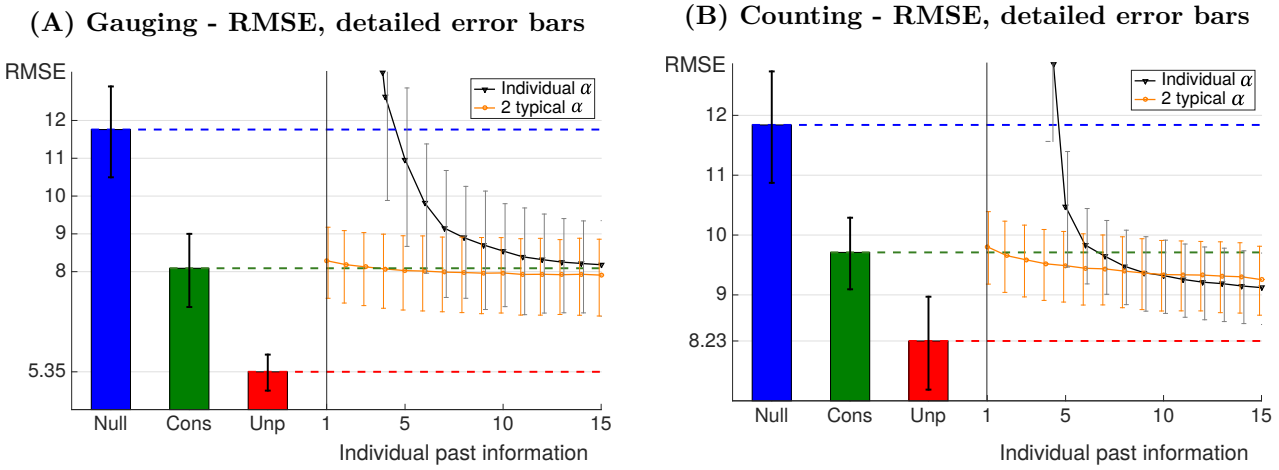

Figure S3. Root mean square error (*RMSE*) of the predictions with detailed error bars for the final round.

S5.2 Prediction accuracy in terms of Mean Absolute Errors

Measuring prediction accuracy in terms of MAEs may appear more intuitive for comparing prediction methods. Fig. S4. assesses the models using an absolute linear scale, where the errors are deliberately unscaled for the *counting game*. The prediction methods rank equally when measured in terms of MAE or RMSE. Notice that, due to nonlinear relation between RMSE and MAE, on this alternative scale, the consensus models errors are now closer from the null model than from the unpredictability error. For comparison, recall that for the *gauging game*, the judgments range between 0 and 100 while they range between 0 and 500 for the *counting game*.

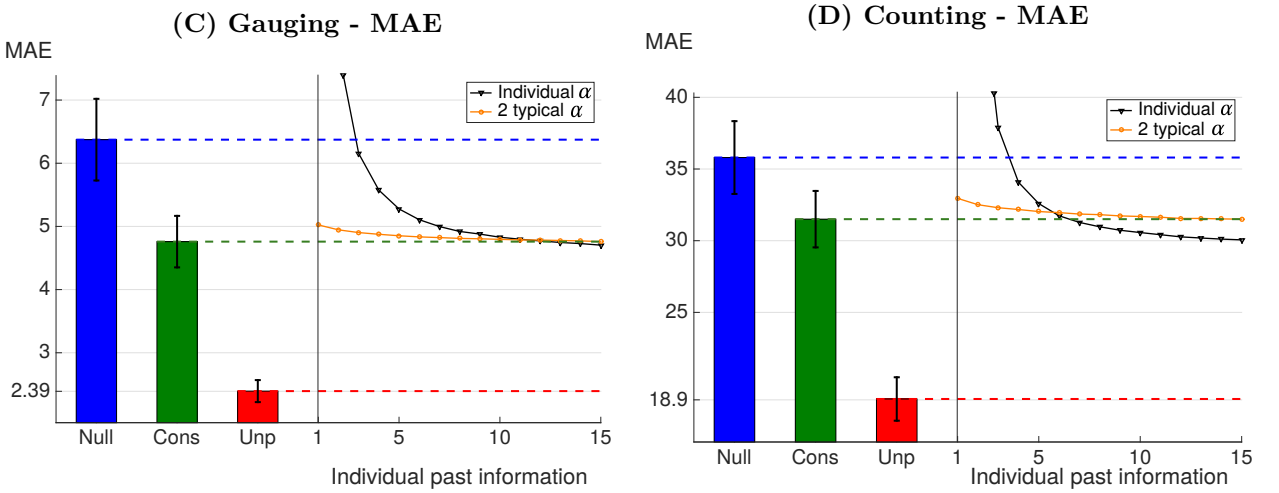

Figure S4. Mean absolute error (*MAE*) of the predictions (unscaled) for the final round.

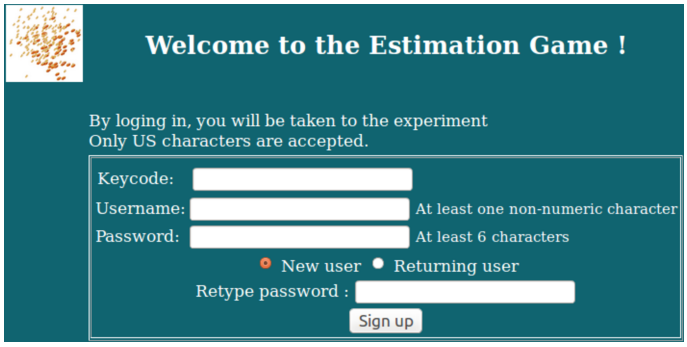

Figure S5. Login page.

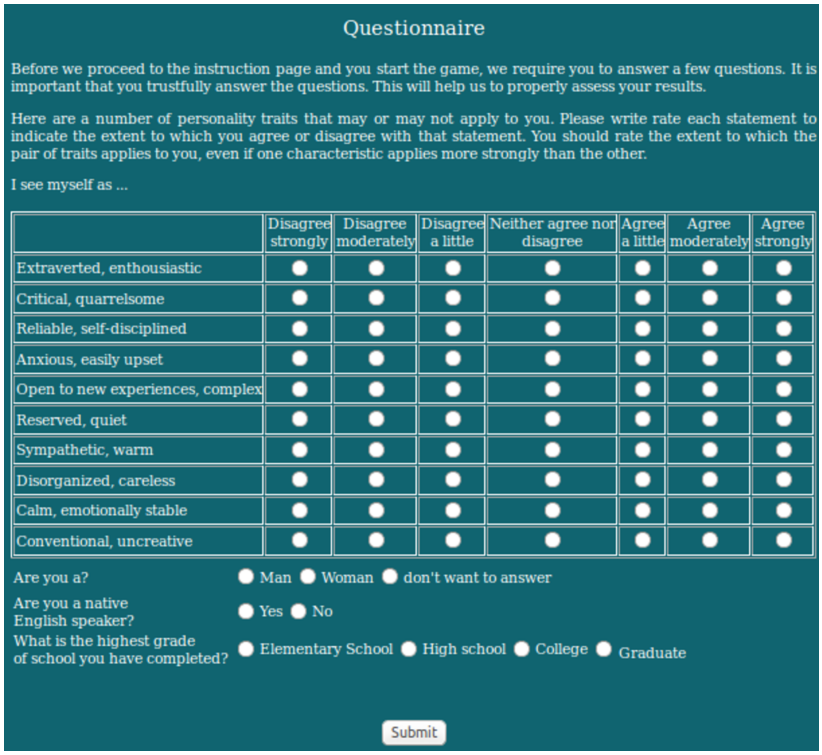
Questionnaire
Before we proceed to the instruction page and you start the game, we require you to answer a few questions. It is important that you trustfully answer the questions. This will help us to properly assess your results.
Here are a number of personality traits that may or may not apply to you. Please write rate each statement to indicate the extent to which you agree or disagree with that statement. You should rate the extent to which the pair of traits applies to you, even if one characteristic applies more strongly than the other.
I see myself as ...
Disagree strongly
Disagree moderately
Disagree a little
Neither agree nor disagree
Agree a little
Agree moderately
Agree strongly
Extraverted, enthousiastic
Critical, quarrelsome
Reliable, self-disciplined
Anxious, easily upset
Open to new experiences, complex
Reserved, quiet
Sympathetic, warm
Disorganized, careless
Calm, emotionally stable
Conventional, uncreative
Are you a?
Man
Woman
don't want to answer
Are you a native English speaker?
Yes
No
What is the highest grade of school you have completed?
Elementary School
High school
College
Graduate
Submit

Figure S6. Questionnaire form.

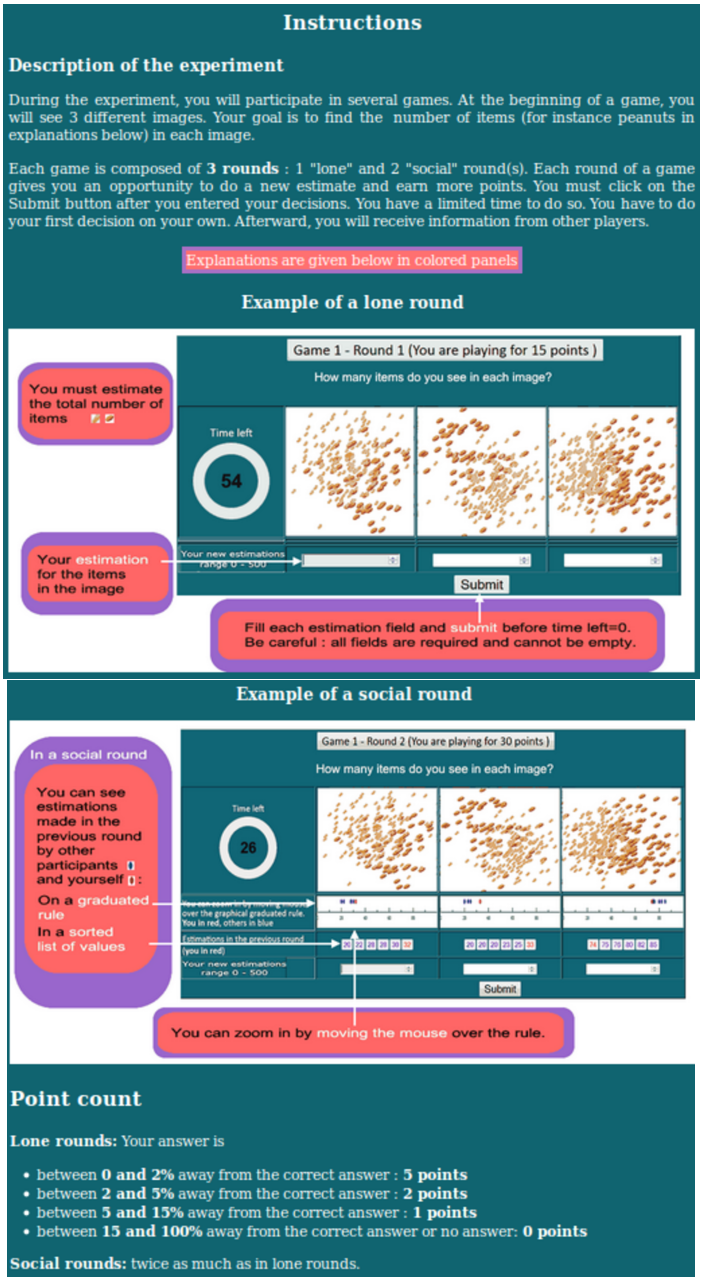
Instructions
Description of the experiment
During the experiment, you will participate in several games. At the beginning of a game, you will see 3 different images. Your goal is to find the number of items (for instance peanuts in explanations below) in each image.
Each game is composed of 3 rounds : 1 "lone" and 2 "social" round(s). Each round of a game gives you an opportunity to do a new estimate and earn more points. You must click on the Submit button after you entered your decisions. You have a limited time to do so. You have to do your first decision on your own. Afterward, you will receive information from other players.
Explanations are given below in colored panels
Example of a lone round
Game 1 - Round 1 (You are playing for 15 points)
How many items do you see in each image?
You must estimate the total number of items
Time left
54
Your estimation for the items in the image
Your new estimations
Submit
Fill each estimation field and submit before time left=0.
Be careful : all fields are required and cannot be empty.
Example of a social round
Game 1 - Round 2 (You are playing for 30 points)
How many items do you see in each image?
In a social round
You can see estimations made in the previous round by other participants and yourself :
On a graduated rule
In a sorted list of values
Time left
26
Estimations in the previous round (you in red)
Your new estimations
range 0 - 500
Submit
You can zoom in by moving the mouse over the rule.
Point count
Lone rounds: Your answer is
between 0 and 2% away from the correct answer : 5 points
between 2 and 5% away from the correct answer : 2 points
between 5 and 15% away from the correct answer : 1 points
between 15 and 100% away from the correct answer or no answer: 0 points
Social rounds: twice as much as in lone rounds.

Figure S7. Instruction page.

(A)

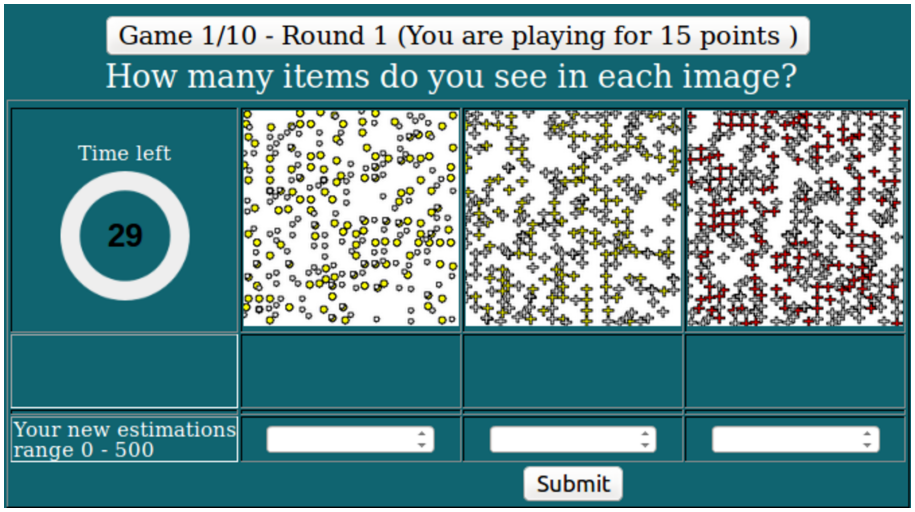
Game 1/10 - Round 1 (You are playing for 15 points)
How many items do you see in each image?
Time left
29
Your new estimations
range 0 - 500
Submit

(B)

Figure S8. (A) Interface of the first round for the *counting game*, played alone. (B) Instances of pictures for the *gauging game*.

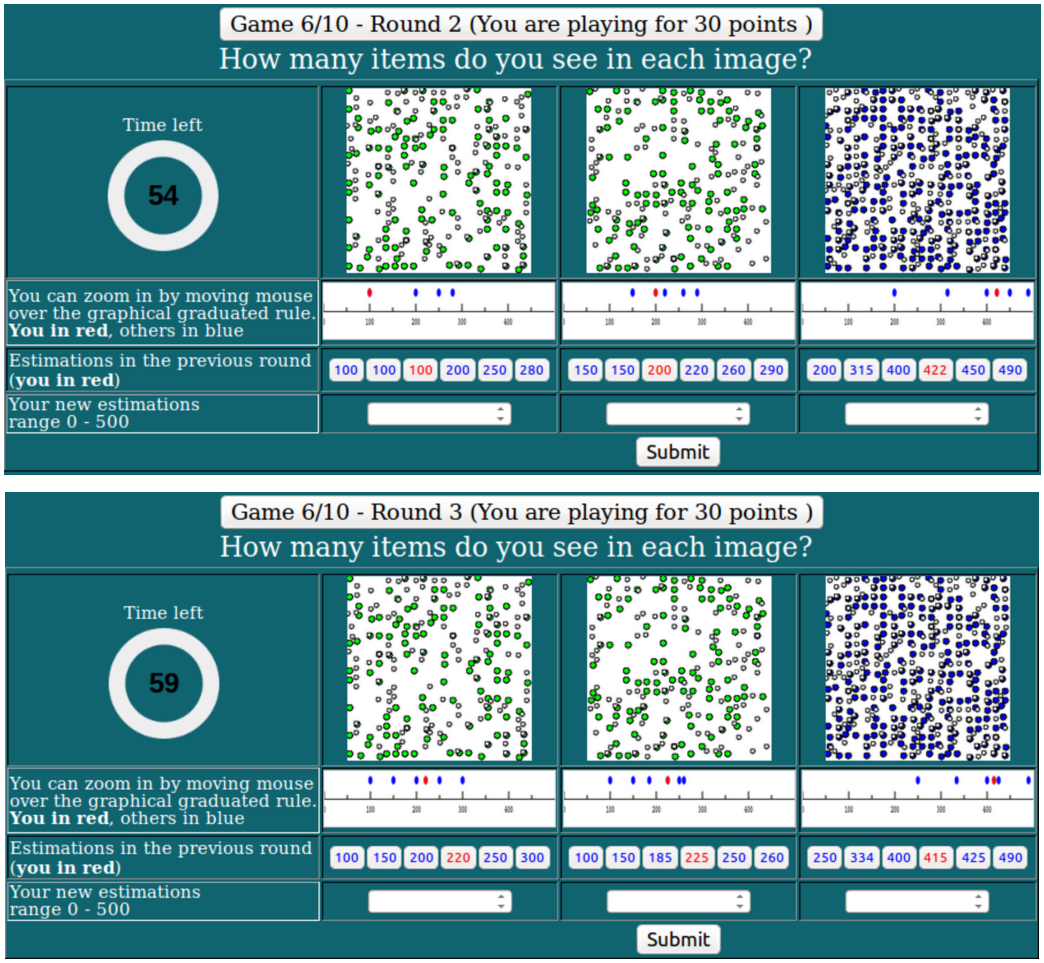
Game 6/10 - Round 2 (You are playing for 30 points)
How many items do you see in each image?
Time left
54
You can zoom in by moving mouse over the graphical graduated rule. You in red, others in blue
Estimations in the previous round (you in red)
100 100 100 200 250 280
150 150 200 220 260 290
200 315 400 422 450 490
Your new estimations range 0 - 500
Submit
Game 6/10 - Round 3 (You are playing for 30 points)
How many items do you see in each image?
Time left
59
You can zoom in by moving mouse over the graphical graduated rule. You in red, others in blue
Estimations in the previous round (you in red)
100 150 200 220 250 300
100 150 185 225 250 260
250 334 400 415 425 490
Your new estimations range 0 - 500
Submit

Figure S9. Interface of the second and third rounds for the *counting game*, social rounds

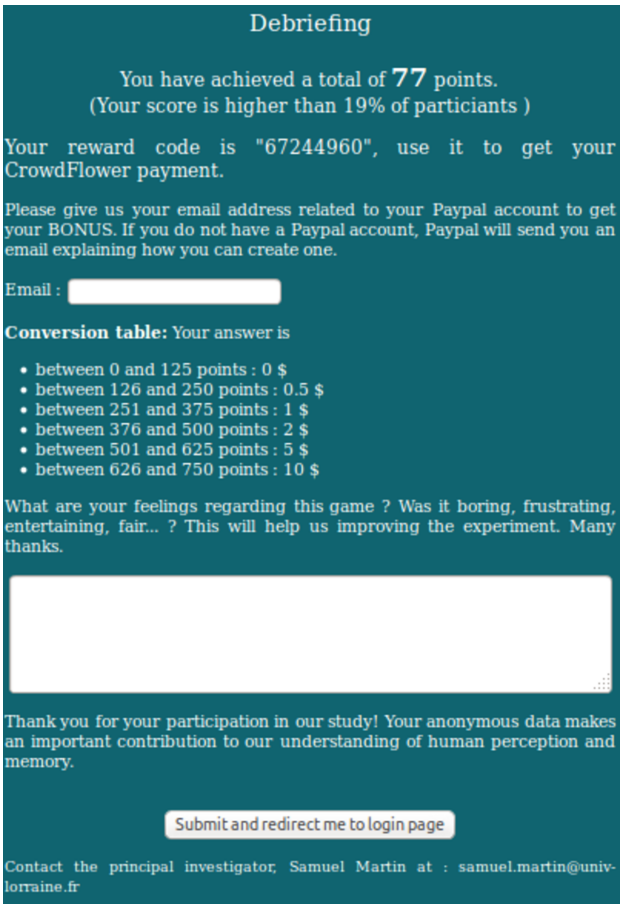
Debriefing
You have achieved a total of 77 points.
(Your score is higher than 19% of particiants)
Your reward code is "67244960", use it to get your CrowdFlower payment.
Please give us your email address related to your Paypal account to get your BONUS. If you do not have a Paypal account, Paypal will send you an email explaining how you can create one.
Email :
Conversion table: Your answer is
between 0 and 125 points : 0 $
between 126 and 250 points : 0.5 $
between 251 and 375 points : 1 $
between 376 and 500 points : 2 $
between 501 and 625 points : 5 $
between 626 and 750 points : 10 $
What are your feelings regarding this game ? Was it boring, frustrating, entertaining, fair... ? This will help us improving the experiment. Many thanks.
Thank you for your participation in our study! Your anonymous data makes an important contribution to our understanding of human perception and memory.
Submit and redirect me to login page
Contact the principal investigator, Samuel Martin at : samuel.martin@univ-lorraine.fr

Figure S10. Debrief page.